Cisco Internet Streamer CDS 2.0–2.3 Software Configuration Guide

April 2009

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-13493-04

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0807R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco Internet Streamer CDS 2.0-2.3 Software Configuration Guide
© 2008 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface xiii

Document Revision History xiii

Audience xiv

Objective xiv

Document Organization xiv

Document Conventions xv

Related Publications xvi

Obtaining Documentation and Submitting a Service Request xvii

CHAPTER 1

Product Overview 1-1

Overview 1-1

Ingest and Distribution 1-3

Prefetch Ingest 1-3

Dynamic Ingest 1-3

Hybrid Ingest 1-3

Live Stream Ingest and Split 1-4

Delivery 1-4

3-Screen Session Shifting 1-5

Management 1-7

Content Delivery System Architecture 1-7

Service Engine 1-7

Content Acquirer 1-8

Internet Streamer 1-8

Service Router 1-20

IP-Based Redirection 1-20

Service Router Workflow 1-21

Coverage Zone File 1-21

Routing Methods 1-22

Request Redirection 1-25

Content Delivery System Manager 1-25

Authentication, Authorization, and Accounting 1-25

Device Management 1-26

Delivery Services Management 1-26

Resiliency and Redundancy 1-27

Content Acquirer Redundancy 1-27
Internet Streamer Redundancy 1-27
Service Router Redundancy 1-27
Internet Streaming CDSM Redundancy 1-27

CHAPTER 2 **Network Design** 2-1

Cisco CDS Topology 2-1
Device Groups 2-2
Baseline Groups 2-3
Delivery Service 2-3
Origin Servers 2-5
Proxy Server 2-5
Manifest File 2-5
Crawling 2-5
Content Acquirer 2-6
SMB Servers 2-6
HTTP Servers 2-7
Internet Streamer 2-7
HTTP Download—Disabling 2-7
Service Workflow 2-7
Programs 2-9
Live Programs 2-9
Rebroadcasts 2-10
API Program File 2-10

CHAPTER 3 **Getting Started** 3-1

Initially Configuring the Devices 3-1
Logging In to the Internet Streaming CDSM 3-1
Activating and Synchronizing the Devices 3-3
Activating and Setting NTP for Each Device 3-3
Activating All Inactive Service Engines 3-5
Navigating the Internet Streaming CDSM 3-6
Devices, Services, and Other Tables 3-7
Device Home Page 3-7
Task Bar 3-8
Configuring Primary and Standby CDSMs 3-8
Changing a Standby to a Primary CDSM 3-9
Typical Configuration Workflow 3-10

CHAPTER 4 **Configuring Devices** 4-1

Configuring Locations 4-1
Configuring Device Groups 4-3
Device Group Overlap 4-5
Configuring the Service Engine 4-6
Activating a Service Engine 4-6
Assigning Devices to Device Groups 4-9
Configuring Bandwidth for Replication and Ingest 4-10
Default Bandwidth 4-11
Scheduled Bandwidth 4-12
Service Control 4-14
Configuring Service Rules 4-15
Configuring ICAP 4-19
Configuring PCMM QoS Policy 4-22
Configuring URL Signing 4-23
Configuring Transaction Logs 4-26
Application Control 4-28
Configuring Default and Maximum Bandwidth 4-29
Configuring Bandwidth Schedules 4-30
Configuring Windows Media Streaming—General Settings 4-32
Configuring Windows Media Streaming Multicast Station 4-36
Configuring Windows Media Streaming Multicast Station Schedule 4-38
Configuring Windows Media Streaming—Broadcast Alias 4-38
Configuring Windows Media Streaming—Bypass List 4-39
Configuring Movie Streamer—General Settings 4-40
Configuring RTSP Advanced Settings 4-42
Configuring Flash Media Streaming 4-42
Configuring Web Engine HTTP Connections 4-43
Configuring Web Engine HTTP Caching 4-45
Configuring Web Engine HTTP Cache Freshness 4-46
Configuring Web Engine Advanced HTTP Caching 4-47
General Settings 4-48
Configuring Content Management 4-48
Login Access Control 4-49
Authentication 4-53
Scheduling Database Maintenance 4-55
Setting Storage Handling 4-56
Network Settings 4-57
Configuring Notification and Tracking 4-76

Configuring Troubleshooting 4-86
Configuring Service Router Settings 4-86
Configuring the Service Router 4-87
Activating a Service Router 4-87
Configuring Routing Settings 4-88
IP-Based Redirection 4-90
Configuring Application Control 4-90
General Settings—Last Resort 4-91
Configuring the CDSM 4-91

CHAPTER 5 **Configuring Services** 5-1
Configuring Delivery Services 5-1
Content Origins 5-1
Creating Delivery Service 5-4
Identifying Content 5-11
Identifying Content Using the CDSM 5-11
Identifying Content Using a Manifest File 5-22
Configuring Programs 5-26
Defining a Program 5-27
Configuring Live Programs 5-27
Priming a Live Delivery Service 5-34
Configuring a Rebroadcast 5-34
Configuring a Multicast Address Pool 5-38
Viewing Programs 5-39
Viewing and Modifying API Programs 5-40
Previewing a Program 5-41
Copying a Program 5-41

CHAPTER 6 **Configuring the System** 6-1
Configuring AAA 6-1
Creating, Editing, and Deleting Users 6-2
Creating, Editing, and Deleting Roles 6-4
Creating, Editing, and Deleting Domains 6-4
Changing a Password 6-5
Configuring System Settings 6-6
System Properties 6-6
Configuring Fast SE Offline Detection 6-7
Configuring Distribution QoS 6-8

Configuring Service Routing 6-9
Coverage Zone File Registration 6-9
Configuring Global Routing 6-11

CHAPTER 7 **Monitoring the Internet Streamer CDS 7-1**
System Monitoring 7-1
System Status 7-1
Device Alarms 7-2
Troubleshooting Devices Using the System Status Bar 7-3
Service Alarms 7-3
System Home Page 7-4
System Audit Logs 7-5
Device Monitoring 7-6
Devices Table 7-6
Device Home Page 7-7
Using show and clear Commands 7-9
Using the CDSM show or clear Command Tool 7-9
CPU Utilization 7-13
Reports 7-14
Bytes Served 7-15
Bandwidth Efficiency Gain 7-15
Streaming Sessions 7-16
Delivery Service Monitoring 7-17
Delivery Services Table 7-17
Replication Status for a Delivery Service 7-20
Content Replication Status by Delivery Service 7-23
Content Replication Status by Device 7-25
Viewing Statistics 7-27
Viewing Service Engines and Device Group Statistics 7-27
Viewing Routing Statistics 7-28
Replication Status 7-29
Transaction Logs 7-31
Transaction Log Formats 7-31
Extended Squid 7-31
Apache 7-32
Custom Format 7-32
Transaction Logging and NTLM Authentication 7-33
Usage Guidelines for Log Files 7-33
Working Logs 7-33

Archive Working Log 7-34
Exporting Log Files 7-34
Windows Media Transaction Logging 7-35

CHAPTER 8 **Maintaining the Internet Streamer CDS** 8-1

Software Upgrade 8-1
Getting a Software File from Cisco.com 8-1
Pre-positioning a Software File 8-2
Finding the Software Version of the Devices 8-3
Configuring the Software Image Settings 8-3
Upgrading the Software 8-6
Interoperability Considerations 8-6
Upgrading Software by Device Groups 8-7
Software Upgrades by Device 8-9
Rebooting Devices 8-9
Deleting a Device 8-10
Deleting a Warm Standby CDSM 8-12
Replacing a Device 8-12
Replacing a CDSM 8-13
Replacing an SE or SR 8-13
Backup and Recovery Procedures 8-15
Performing Backup and Restore on the CDSM Database 8-15
Using the Cisco CDS Software Recovery CD-ROM 8-16
System Software Components 8-16
Getting the Cisco CDS Software Recovery File from Cisco.com 8-17
Installing the Software Using the Recovery CD-ROM 8-17
Recovery CD-ROM Options 8-18
Recovering the System Software 8-19
Recovering a Lost Administrator Password 8-21
Recovering from Missing Disk-Based Software 8-22
Recovering CDS Network Device Registration Information 8-25
Disk Maintenance 8-26
Removing and Replacing Hard Disk Drives 8-26
Removing and Replacing a cdnfs Disk 8-26
Removing and Replacing a sysfs Disk 8-27
Removing Content 8-28
Single Content Removal 8-28
Batch Content Removal 8-28

APPENDIX A **Troubleshooting** A-1

Troubleshooting Service Router Configurations A-1
Troubleshooting the Distribution Hierarchy A-2
Troubleshooting Content Acquisition A-3
Enabling the Kernel Debugger A-6

APPENDIX B **Creating Manifest Files** B-1

Introduction B-1
Manifest File Requirements B-2
Working with Manifest Files B-2
Specifying a Single Content Item B-2
Specifying a Crawl Job B-3
Writing Common Regular Expressions B-5
Scheduling Content Acquisition B-5
Specifying Shared Attributes B-6
Specifying a Crawler Filter B-7
Specifying Content Priority B-8
Generating a Playserver List B-9
Customized Manifest Playserver Tables and the HTTP Playserver B-10
Specifying Attributes for Content Serving B-10
Specifying Time Values in the Manifest File B-11
Refreshing and Removing Content B-12
Specifying Live Content B-13
Specifying Hybrid Ingest Content B-13
Manifest Validator Utility B-14
Running the Manifest Validator Utility B-14
Valid Manifest File Example B-14
Invalid Manifest File Example B-15
Understanding Manifest File Validator Output B-16
Syntax Errors B-17
Syntax Warnings B-17
Correcting Manifest File Syntax B-17
Manifest File Structure and Syntax B-17
CdnManifest B-21
playServerTable B-22
playServer B-22
options B-24
server B-25

host B-25
proxyServer B-27
item B-28
crawler B-35
item-group B-38
matchRule B-41
match B-41
contains B-43
XML Schema B-44
Manifest XML Schema B-44
PlayServerTable XML Schema B-48
Default PlayServerTable Schema B-49
Manifest File Time Zone Tables B-50

APPENDIX C Creating Coverage Zone Files C-1
Introduction C-1
Coverage Zone XML Schema C-2
Coverage Zone File Examples C-3
Scenario 1: Coverage Zone with Client Network Only C-3
Scenario 2: Coverage Zone with Geographical Location of the Datacenter Only C-4
Scenario 3: Coverage Zone with Client Network and Geographical Location of the Datacenter C-4

APPENDIX D Creating and Manipulating Session Shifting Files D-1
Introduction D-1
Session Shifting and Delivery Services D-2
Content Manager D-2
Content Lists D-2
Supported Content Management Operations D-3
Subscriber Manager D-4
Subscriber Lists D-4
Supported Subscriber Management Operations D-5
Profile Manager D-5
Profile Lists D-6
Profiles in Content Lists D-6
Supported Profile Management Operations D-7
Session Manager D-7
Stream Session Event D-8
Session Update D-8

Session Lookup D-9

APPENDIX E **CLI Commands** E-1

Configuring Port Channel E-1

Redundant Dedicated Management Ports E-1

Configuring Last-Resort Routing E-6

Configuring Standby Interfaces E-6

Other CLI Commands E-10

CDNFS cleanup Command E-10

Disk Commands E-11

cache content Command E-12

APPENDIX F **URL Signing and Validation** F-1

Introduction F-1

URL Signing Components F-1

Supported Protocols and Media F-2

Configuring the CDS for URL Signing F-2

URL Signing Script F-2

Example of a Python URL Signing Script F-2

Running a Python URL Signing Script F-3

Importance of Device Synchronization F-4

Understanding the Signing Procedure F-5

APPENDIX G **Testing the Internet Streamer CDS** G-1

Testing the Web Engine G-1

Testing Preingested Web Content G-1

Testing Dynamically Ingested Web Content G-3

Testing the Windows Media Streaming Engine G-7

Testing Preingested Windows Media Content G-7

Testing Dynamically Ingested Windows Media Content G-8

Testing Windows Media Live Content Playback G-10

Testing the Movie Streamer Engine G-11

Preparing Movie Streamer Content for Ingest G-12

Testing Preingested Movie Streamer Content G-13

Testing Dynamically Ingested Movie Streamer Content G-16

Testing Movie Streamer Live Content Playback G-17

Testing the Flash Media Streaming Engine G-19

Testing Flash Media Streaming Preingested Content G-20

Testing Flash Media Streaming Dynamically Ingested Content G-22
Testing Flash Media Streaming—Live Streaming G-24

APPENDIX H **Software Licensing Information** H-1

Notices H-1
OpenSSL/Open SSL Project H-1
License Issues H-1

Preface

This preface describes the audience, use, and organization of the *Cisco Internet Streamer CDS 2.0-2.3 Software Configuration Guide*. The preface also outlines the document conventions and support information.

This preface contains the following sections:

- Document Revision History, page xiii
- Audience, page xiv
- Objective, page xiv
- Document Organization, page xiv
- Document Conventions, page xv
- Related Publications, page xvi
- Obtaining Documentation and Submitting a Service Request, page xvii

Document Revision History

The Document Revision History table below records technical changes to this document.

Document Version	Date	Notes
OL-13493-04	October 4, 2008	The CDS Release 2.3 Flash Media Streaming interactive application functionality and Service Router IP-based redirection were added, as well as other enhancements.
OL-13493-03	April 4, 2008	The CDS Release 2.2 service-aware and content-aware functions were added to the Service Router functionality. Dynamic live streaming for Flash Media Streaming, the 3-Screen Session Shifting feature, and information on support for H.264 were added to this document.
OL-13493-02	March 24, 2008	The CDS Release 2.1 Flash Media Streaming feature was added to this document.
OL-13493-01	August 20, 2007	Initial release.

Note All information, fields, and features that are marked as a Release 2.1 feature are also available in Releases 2.2 and 2.3.
All information, fields, and features that are marked as a Release 2.2 feature are also available in Release 2.3.
All information, fields, and features that are marked as a Release 2.0 feature are only available in Release 2.0.

Audience

This guide is for the networking professional managing the Cisco Internet Streamer Content Delivery System, hereafter referred to as the *CDS*. Before using this guide, you should have experience working with the Cisco IOS software and be familiar with the concepts and terminology of Ethernet, local area networking, and Internet streaming.

Objective

This guide provides the information you need to configure and monitor the Cisco Internet Streamer CDS.

This guide provides procedures for using the commands that have been created or changed for use with the Cisco Internet Streamer CDS. It does not provide detailed information about these commands.

This guide does not describe system messages you might encounter or how to install your CDS. For information on installing the hardware, see the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide*. See the "Related Publications" section on page xvi for links to documentation online.

For documentation updates, see the release notes for this release.

Document Organization

This document contains the following chapters and appendices:

Chapter or Appendix	Description
Chapter 1, "Product Overview"	Provides a brief introduction to the Internet Streamer CDS.
Chapter 2, "Network Design"	Describes the CDS topology, elements of a delivery service, and the delivery service workflow.
Chapter 3, "Getting Started"	Provides information about initially configuring the devices to communicate with the CDSM, configuring a standby CDSM, navigating the CDSM, and a typical configuration workflow.
Chapter 4, "Configuring Devices"	Provides information on configuring the devices in the Internet Streamer CDS.
Chapter 5, "Configuring Services"	Provides information about configuring delivery services.

Chapter or Appendix	Description
Chapter 6, "Configuring the System"	Provides information on system configuration for the Internet Streamer CDS.
Chapter 7, "Monitoring the Internet Streamer CDS"	Provides information on monitoring the Internet Streamer CDS.
Chapter 8, "Maintaining the Internet Streamer CDS"	Provides information on upgrading the Internet Streamer CDS software, deleting devices from the system, performing disk maintenance, and removing content from the system.
Appendix A, "Troubleshooting"	Discusses troubleshooting Service Routers, and the acquisition and distribution of content.
Appendix B, "Creating Manifest Files"	Provides information on creating and validating a Manifest file.
Appendix C, "Creating Coverage Zone Files"	Provides information on creating and validating a Coverage Zone file.
Appendix D, "Creating and Manipulating Session Shifting Files"	Provides information on creating Session Shifting XML files.
Appendix E, "CLI Commands"	Provides information on configuring port channels, last resort routing domains, and other CLI commands.
Appendix F, "URL Signing and Validation"	Describes the URL signing script for URL signature creation at the portal.
Appendix G, "Testing the Internet Streamer CDS"	Describes procedures for testing the Internet Streamer CDS using the different media players.
Appendix H, "Software Licensing Information"	Provides information about the software licenses for the Internet Streamer CDS.

Document Conventions

This guide uses the following conventions for command syntax descriptions and textual emphasis:

Convention	Description
boldface font	Commands and keywords are in **boldface**.
italic font	Arguments for which you supply values are in *italics*.
[ ]	Elements in square brackets are optional.
{*x* | *y* | *z*}	Alternative, mutually exclusive, keywords are grouped in braces and separated by vertical bars.
[*x* | *y* | *z*]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
`screen` font	Terminal sessions and information the system displays are in `screen` font.
`boldface screen` font	Information you must enter is in **`boldface screen`** font.
`italic screen` font	Arguments for which you supply values are in *`italic screen`* font.

Convention	Description
^	The symbol ^ represents the key labeled Control—for example, the key combination ^D in a screen display means hold down the Control key while you press the D key.
< >	Nonprinting characters, such as passwords, are in angle brackets in contexts where italics are not available.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Note Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this publication.

Tip Means the following information might help you solve a problem.

Related Publications

These documents provide complete information about the CDS and are available from the Cisco.com site:

- *Release Notes for Cisco Internet Streamer CDS 2.0*
- *Release Notes for Cisco Internet Streamer CDS 2.1*
- *Release Notes for Cisco Internet Streamer CDS 2.2*
- *Release Notes for Cisco Internet Streamer CDS 2.3*
- *Cisco Internet Streamer CDS 2.0-2.3 Quick Start Guide*
- *Cisco Internet Streamer CDS 2.0-2.3 API Guide*
- *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide*
- *Cisco Content Delivery System 2.x Documentation Roadmap*
- *Regulatory Compliance and Safety Information for Cisco Content Delivery Engine 100/200/300/400*

You can access the software documents at the following URL:

http://www.cisco.com/en/US/products/ps7127/tsd_products_support_series_home.html

You can access the hardware documents at the following URL:

http://www.cisco.com/en/US/products/ps7126/tsd_products_support_series_home.html

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

CHAPTER 1

Product Overview

This chapter provides an introduction to the Cisco Internet Streamer Content Delivery System (CDS). This chapter has the following major topics:

- Overview, page 1-1
- Content Delivery System Architecture, page 1-7

Overview

The Cisco Content Delivery System (CDS) is a distributed network of Content Delivery Engines (CDEs) running Content Delivery Applications (CDAs) that collaborate with each other to deliver multi-format content to a variety of client devices. The client devices supported in Releases 2.0–2.3 are personal computers and Wi-Fi-enabled mobile devices, such as personal digital assistants (PDAs).

The CDS supports a variety of mechanisms to accelerate the distribution of content within the content delivery network. The CDS offers an end-to-end solution for service providers to ingest and stream entertainment-grade content to subscribers.

The CDS functionality can be separated into four areas:

- Ingest
- Distribution
- Delivery
- Management

Each CDE in the CDS contributes to one or more of these functions as determined by the CDAs running on it. Table 1-1 describes the relationship between the CDA names and the Internet Streaming Content Delivery System Manager (CDSM) device names.

Table 1-1 CDA Mapping to Functionality and CDSM

CDA Name	Functionality	CDSM Device Name
Internet Streamer (+ Content Acquirer)	Ingest, distribution, and delivery	Service Engine (SE)
Service Router	Redirect client requests for delivery	Service Router (SR)
Internet Streaming Content Delivery System Manager	Management	CDSM

The Service Engine can function as a Content Acquirer and Internet Streamer, or just as an Internet Streamer.

Figure 1-1 shows the major elements of a CDS network. How content flows, from ingest to distribution within the CDS, to delivery to client devices, is dictated by the content delivery services defined for each Content Origin. A delivery service is a configuration defined by using the CDSM and consists of configuration parameters that dictate how content is ingested and distributed, and what content is delivered to the client devices. Some of the primary delivery service definition parameters are:

- Origin server
- Service routing domain name
- Service Engines participating in the delivery service
- Service Engine designated as the Content Acquirer

The Content Acquirer is only active on one Service Engine in each delivery service.

Figure 1-1 High-Level View of the Cisco CDS

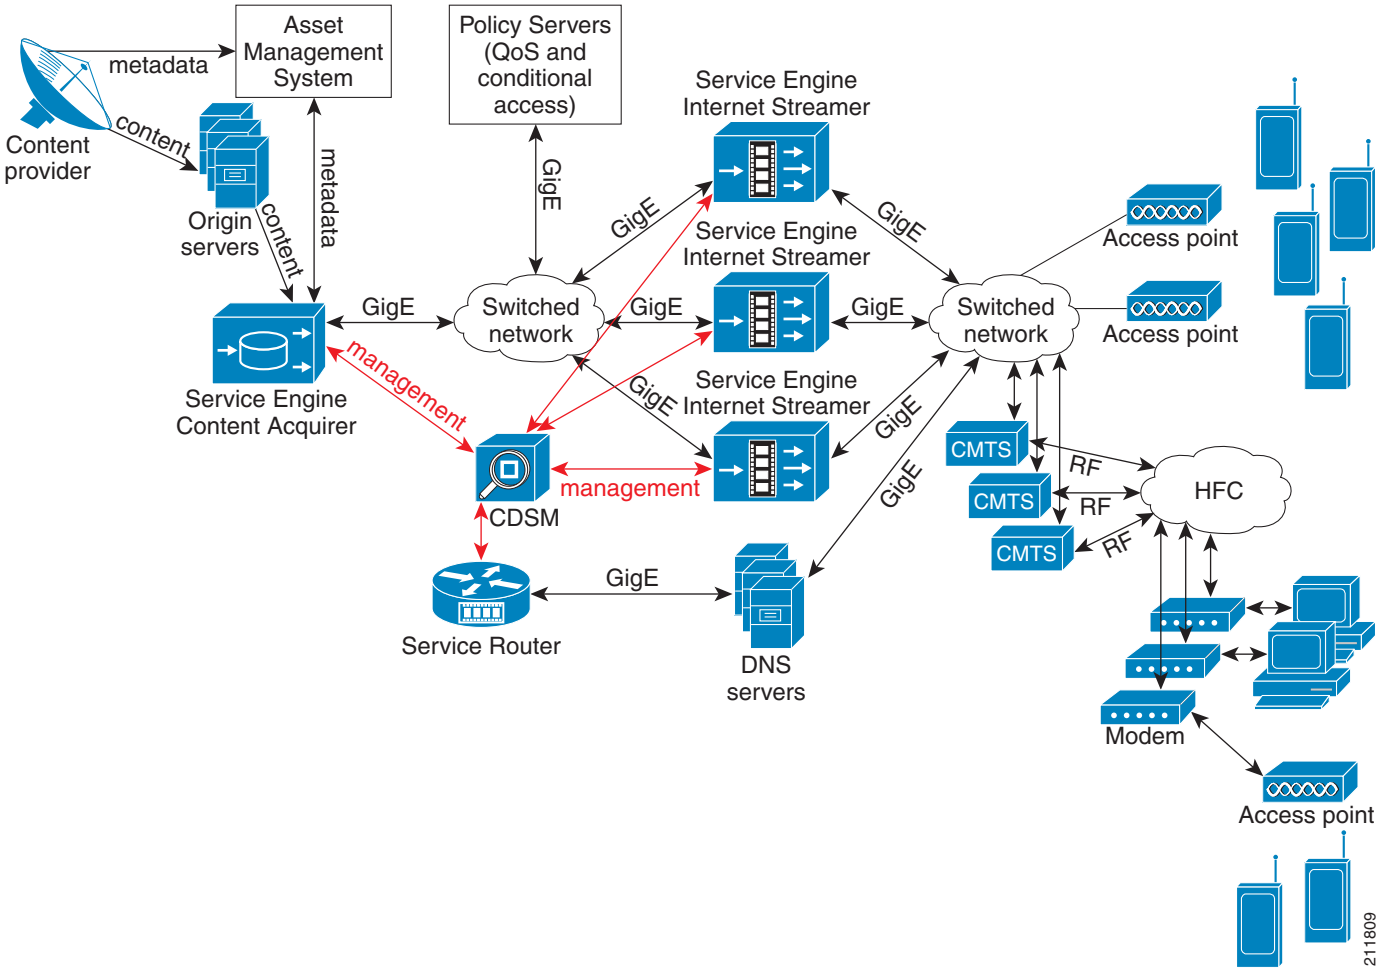

The following sections briefly describe the elements of the CDS. For more detailed information, see the "Content Delivery System Architecture" section on page 1-7.

Ingest and Distribution

The Service Engine designated as the Content Acquirer for a delivery service is the ingest device. Cisco Internet Streamer CDS Releases 2.0–2.3 supports the following methods of content ingest:

- Prefetch ingest
- Dynamic ingest
- Hybrid ingest
- Live stream ingest and split

The distribution of content within the CDS is determined by the method of ingest used.

Note The recommended maximum number of prefetched content items is 200,000.

Prefetch Ingest

The Content Acquirer receives metadata from the backoffice in the form of an XML-formatted Manifest file and, using the information in the file, pulls the content into storage on the Content Acquirer. The content can be ingested by using different protocols. The supported protocols are FTP, HTTP, HTTPS, CIFS, as well as local files, which are files copied to the Service Engine. The ingested content is then distributed to all Service Engines in the content delivery service. The content is stored on each Service Engine's hard disk for a configurable amount of time or until the content entry gets deleted from the Manifest file. This is called *content pinning*.

The Manifest file can be used to specify different policies for content ingest and also for streaming the prefetched content. For example, the policy could include specifying the expiry of the content, setting time windows in which the content is made available to users, and so on.

Dynamic Ingest

Content can be dynamically ingested into the CDS. Dynamic ingest is triggered when a Service Engine's Internet Streamer application does not find a client's requested content in its local hard disk storage. All Service Engines participating in the content delivery service coordinate to form a content distribution tunnel starting at the origin server and ending at the Service Engine responding to the client request. As the content flows through this tunnel, the participating Service Engines cache a copy of the content. Subsequent requests for the same content are served off the CDS network. Content ingested and distributed by this method is deleted if clients do not request it frequently.

The Internet Streaming CDSM manages this ingest method internally, not by instructions embedded in a Manifest file, and manages the storage automatically. The Internet Streaming CDSM also provides the ability to purge any dynamically ingested content out of the Service Engines. Content is identified by a URL, which is also used to delete the content.

Hybrid Ingest

The hybrid ingest method provides a very powerful solution by combining the features of the prefetch ingest and the dynamic ingest methods. The metadata and control information about the content, defined in the Manifest file, is propagated and pinned to all Service Engines participating in the content delivery service. However, the content is not prefetched. Ingest occurs upon user request for the content. Content

that is cached on the Service Engines by using this method is subject to the same deletion rules as the dynamic ingest method. The metadata that is propagated can be used to specify explicit controls and policies for streaming the content.

Live Stream Ingest and Split

The live stream ingest method distributes a live content feed to all the Service Engines participating in the content delivery service and helps to scale the content delivery to a very large audience. This method leverages the live stream splitting capabilities of the Internet Streamer application and optimizes the access by doing a one-to-many split to all Service Engines in the content delivery service. The Internet Streaming CDSM provides the necessary interface to schedule the streaming of live programs. Advanced techniques are used to enhance the performance of live streaming.

Delivery

The Service Router handles client requests for content and determines the best Service Engine to deliver it based on proximity, load and health states.

Once the best Service Engine has been determined, the content is delivered to the client device by means of one of the following mechanisms:

- **Static Content Download using HTTP**—Content is downloaded by the client device before it can be rendered to the user.
- **Progressive Content Download using HTTP**—Content is rendered in segments to the user before it has been fully downloaded.
- **Content Streaming using HTTP, RTMP, RTSP, or RTP**—Content is streamed to the client device, Service Engines collect feedback and can fine-tune streaming. Advanced error recovery can also be performed. This is a very common method of streaming video content to client devices.

Table 1-2 lists the content types and formats, content transport protocols, and client types supported by the CDS.

Table 1-2 Supported Content Types

Content Types and Formats	Transport Protocols	Typical Client Types	Access Network Type
Windows Media (WMA, WMV, ASF, and others) VC-1	RTP, RTSP, HTTP	Windows Media Player 9, 10, 11 on PC Windows Media Player 9 for Mac Windows Media Player 9 on PDA running Windows Pocket PC 2002/2003 Windows Media Player 10 on PDA running Windows Mobile 5 Windows Media Technology (WMT) Silverlight	Wired Wi-Fi Cellular
QuickTime (MOV), hinted (3GP) files	RTP, RTSP, HTTP	On PC: QuickTime Player, QuickTime Pro 6 or 7, RealPlayer 10 or 11 (3GP only), VLC player On Mac: QuickTime Player, QuickTime Pro 6 or 7, RealPlayer 10 for Mac OS X (3GP only), On Mobile: RealPlayer on Nokia N series phones (3GP only), PackerVideo player	Wired Wi-Fi Cellular

Table 1-2 Supported Content Types (continued)

Content Types and Formats	Transport Protocols	Typical Client Types	Access Network Type
Other Hypertext and image files (HTML, JPEG, and so on)	HTTP	Web browsers and other HTTP clients	Wired Wi-Fi Cellular
MPEG (MP1, MP2, MP4)	RTP, RTSP	MPEG clients **Note** For Flash Media Streaming, the Adobe Flash Media Player 9 update 3, Adobe Media Player, and Adobe Air, are the only players that support MPEG-4.[1]	Wired
Adobe Flash (SWF, FLV, MP3)	RTMP, HTTP	Adobe Flash Player 9 for Windows, Mac OS, and Linux	Wired Wi-Fi Cellular
H.264[2]	RTMP, HTTP	H.264 clients **Note** For Flash Media Streaming, the Adobe Flash Media Player 9 update 3, Adobe Media Player, and Adobe Air, are the only players that support H.264.	Wired

1. Support for MPEG-4 in the Flash Media Streaming Engine is a Release 2.2 feature.
2. Support for H.264 in the Flash Media Streaming Engine is a Release 2.2 feature.

Note RTMP is part of the Flash Media Streaming feature and is part of Releases 2.0–2.3

3-Screen Session Shifting

Note 3-Screen Session Shifting is a Release 2.2 feature and supports RTSP streaming for the Windows Media Engine and Movie Streamer Engine.

The 3-Screen Session Shifting feature unifies the user interactions with different content streaming engines and enables seamless movement of user sessions to and from different client devices (PC, Mac, Mobile, and TV). With this feature, users can pause a streaming session on one device and resume it on a different device. This feature provides the intelligence for client detection and format selection. The 3-Screen Session Shifting interacts with the TV CDS as well.

3-Screen Session Shifting integrates with the delivery service configuration by linking three XML files that interact with content management and transcoding systems to obtain device-type mappings and content format mapping. 3-Screen Session Shifting works with entitlement servers to get the user identification and uses the user identification as a key to store user sessions.

Session shifting from one client device to another device for a particular user requires the subscriber ID, content ID, and last play time.

The Service Router acts as the centralized session manager for session shifting and maintains the session database. The session database is populated by operator-generated 3-screen maps (XML files), consisting of the Content map, Subscriber map, and Profile map. The XML schema for these different

XML files is available upon request. The operator can enable session shifting on a delivery service basis and also configure a TV streamer IP address and port for each delivery service. The operator also needs to specify which Service Router has the session database.

The Content map aggregates same semantic content across the TV CDS and Internet Streamer CDS systems. Similarly, the Subscriber map helps to correlate different identifiers of the same users between the TV CDS and Internet Streamer CDS.

Note The Fast Cache feature for the Windows Media Engine is not supported in session shifting.

3-Screen Web Services Infrastructure

Note Release 2.2 supports web-based programming APIs. These APIs are based on the REpresentational State Transfer (REST) architecture. For more information, see Appendix D, "Creating and Manipulating Session Shifting Files." In Release 2.2, 3-Screen Session Shifting is the only application that is supported through the Web Services infrastructure.

Following are the Web Services infrastructure features that are supported in this release:

- Content Manager
- Subscriber Manager
- Profile Manager
- Stream Session Event Manager

Content Manager

The Content Manager facilitates the management of 3-screen content for session shifting. Content lists can be posted to the Web Services interface for content addition or deletion. The provisioned content is visible by using a web browser and entering a simple URL. The schema for the ContentList XML file is available upon request.

Subscriber Manager

The Subscriber Manager facilitates the management of 3-screen users for session shifting. Subscriber lists can be posted to the Web Services interface to add or delete subscribers. The provisioned subscribers are visible by using a web browser and entering a simple URL. The schema for the SubscriberList XML file is available upon request.

Profile Manager

The Profile Manager facilitates the management of different client devices and player types for session shifting. This is called *Capability Exchange*. Capability Exchange refers to the process of mapping the logical URL to a specific content format inside the group. One profile groups together multiple clients that have similar characteristics, such as supported file format, and desired resolution or bandwidth. Profile lists can be posted to the Web Services interface for addition or deletion. The profiles are visible by using a web browser and entering a simple URL. The schema for the ProfileList XML file is available upon request.

Stream Session Event Manager

The Stream Session Event Manager allows creation and retrieval of last play times for specific pairs of content and subscribers. The schema for the StreamSessionEvent XML file is available upon request.

Management

The Internet Streaming CDSM, a secure Web browser-based user interface, is a centralized system management device that allows an administrator to manage and monitor the entire CDS network. All devices, Service Engines and Service Routers, in the CDS are registered to the Internet Streaming CDSM.

Service Engines can be organized into user-defined device groups to allow administrators to apply configuration changes and perform other group operations on multiple devices simultaneously. One device may belong to multiple device groups.

The Internet Streaming CDSM also provides an automated workflow to apply a software image upgrade to a device group.

Content Delivery System Architecture

The CDS consists of an Internet Streaming CDSM, one or more Service Engines, and one Service Router. For full redundancy, a CDS would include an additional CDSM and Service Router. The Service Engine handles content ingest, content distribution within the CDS, and content delivery to client devices. The Service Router handles client requests and redirects the client to the most appropriate Service Engine. The Internet Streaming CDSM manages and monitors the CDS, the delivery services, and all the devices in the CDS.

This section covers the following topics:

- Service Engine
- Service Router
- Content Delivery System Manager
- Resiliency and Redundancy

Service Engine

Each Service Engine can function as a Content Acquirer and Internet Streamer, or just as an Internet Streamer. Based on the Service Engines' assignments to different delivery services, the right set of applications supporting the functions is enabled. For example, only one Service Engine is assigned the role of Content Acquirer in each delivery service. In addition, the Service Engine assigned as the Content Acquirer in a delivery service also includes the functions of an Internet Streamer.

Both the Content Acquirer and the Internet Streamer applications have storage and distribution functions within the CDS, which include the following:

- Management of the physical storage of content and metadata. Content URLs are translated into their physical file paths for content retrieval, deletion, and update.
- Management of dynamically ingested content and periodic replacement of content not accessed frequently. Content replacement is performed by means of sophisticated content-replacement algorithms. The algorithms add "weight" to the content according to size, frequency of access, and other attributes to produce the list of content that needs to be purged.
- Ingest of prefetched content and retrieval of such content for distribution to other Service Engines in the same delivery service.

- Maintenance of information about the entire CDS topology and all the delivery services. This includes upkeep of a list of Service Engines in the same delivery service that is used for distributing prefetched, dynamic, and live stream content.
- Maintenance of the database that stores and distributes metadata about the content, and the topology and delivery service information.
- Distribution of content on a per-delivery service basis, where the flow path of content could differ from one delivery service to another.

Content Acquirer

Every delivery service requires a Content Acquirer, which is a CDA that resides on every Service Engine. The Content Acquirer CDA becomes active when the Service Engine is designated as the Content Acquirer in a delivery service. The Content Acquirer has the following functions and capabilities:

- Fetches content from origin servers using HTTP, HTTPS, FTP, or CIFS (Dynamic ingest supports HTTP only).
- Supports the NT LAN Manager (NTLM) and basic authentication for ingesting content from the origin servers.
- Creates and distributes the metadata for each of the prefetched contents according to the Manifest file and the information returned by the origin server.

Once the Content Acquirer has ingested the content and distributed the metadata, it creates a database record for the metadata and marks the content ready for distribution. All other types of ingest (dynamic, hybrid, and live stream) are handled by the Content Acquirer as well.

Internet Streamer

All Internet Streamers participating in a delivery service pull the metadata from a peer Internet Streamer called a *forwarder*, which is selected by the internal routing module. Each Internet Streamer participating in a delivery service has a forwarder Internet Streamer. The Content Acquirer is the top-most forwarder in the distribution hierarchy. In the case of prefetched ingest, each Internet Streamer in the delivery service looks up the metadata record and fetches the content from its forwarder. For live or cached content metadata, only the metadata is distributed.

The content associated with the metadata for live and cached content is fetched by the specified protocol engine, which uses the dynamic ingest mechanism. When a request for a non-prefetched content arrives at an Internet Streamer, the protocol engine application gets the information about the set of upstream Internet Streamers through which the content can be acquired. In the case of dynamic ingest, the Internet Streamer uses the cache routing function to organize itself as a hierarchy of caching proxies and performs a native protocol cache fill. Live stream splitting is used to organize the Internet Streamers into a live streaming hierarchy to split a single incoming live stream to multiple clients. The live stream can originate from external servers or from ingested content. Windows Media Engine, Movie Streamer Engine, and Flash Media Streaming Engine support live stream splitting.

The Internet Streamers use service control to filter and control incoming requests for content. The service rules and the PacketCable Multimedia (PCMM) Quality of Service (QoS) control are some of the functions that are encapsulated under the Service Control option in the Internet Streaming CDSM.

The Internet Streamers send keep-alive and load information to the Service Router that is participating in the same delivery service. This information is used by the Service Router to choose the most appropriate Internet Streamer to handle the request.

The Internet Streamer function is implemented as a set of protocol engine applications. The protocol engine applications are:

- Web Engine
- Windows Media Engine
- Movie Streamer Engine
- Flash Media Streaming Engine

Web Engine

All HTTP client requests that are redirected to a Service Engine by the Service Router are handled by the Web Engine. On receiving the request, the Web Engine uses its best judgment and either handles the request or forwards it to another component within the Service Engine. The Web Engine, using HTTP, can serve the request from locally stored content in the CDS or from any upstream proxy or origin server.

An HTTP client request that reaches the Service Engine can either be from a Service Router redirect or from a direct proxy request.

On receiving an HTTP request for content, the Web Engine decides whether the content needs to be streamed by the Windows Media Engine, and if so, hands the request over to the Windows Media Engine, otherwise the request is handled by the Web Engine.

Policy Server Integration

The policy control server, a third-party PCMM-compliant system, ensures guaranteed bandwidth for multimedia data delivered over broadband networks. The Web Engine communicates with the policy server by means of Internet Content Adaptation Protocol (ICAP) and HTTP to set and monitor QoS attributes for each client session and whether access should be denied.

Using ICAP, the Web Engine determines whether the bandwidth reservation can be allocated for this client or access should be denied. The policy server uses the cookie generated by the web portal and the client's IP address to make the decision and replies accordingly to the Web Engine.

Using HTTP, the Web Engine communicates the start and teardown of the request to the policy server. Bandwidth reservation is performed when the download starts, and once the download is complete, the Web Engine sends a teardown message to the policy server.

The Web Engine uses PCMM to interact with the policy server to allocate guaranteed bandwidth for authenticated client requests for content. The PCMM integration allows the granting of QoS for the session as well as conditional access protection of the content.

Upon receiving permission from the policy server, the Web Engine generates a URL signature and appends it to the requested URL; it then embeds the new URL in an .asx file and sends the file back to the client. The file consists of the signed URL with RTSP and HTTP options. For RTSP and HTTP streaming, the Windows Media Engine communicates with the policy server for bandwidth commitment. If RTSP and HTTP streaming fail, the client device begins HTTP progressive download of the file. The Web Engine handles the QoS for HTTP progressive download requests.

The signed URL adds additional security. The URL signature generation is based on a key that is a shared secret between the component generating the URL signature and the component validating the URL signature. The URL signature can be generated by the Web Engine, another component external to the Service Engine, or the web portal.

Cache Fill Operations

The Web Engine interfaces with the storage function in the Service Engine to determine whether the content is present locally or whether the content needs to be fetched from either an upstream Service Engine or the origin server.

The Web Engine communicates to the upstream Service Engine for cache fill operations. This interaction is based on HTTP. This cache fill operation is on demand and hence only occurs when the content is not stored locally. The upstream Service Engine can be selected dynamically by means of the Hierarchical Cache Routing Module, or can be configured statically through the Internet Streaming CDSM. The Hierarchical Cache Router generates a list of upstream Service Engines that are alive, ready to serve the request, and part of the delivery service. If the Web Engine is unsuccessful in locating the content on one of these Service Engines, the content is retrieved from the origin server.

Whether the content is found locally or retrieved and stored through the cache fill operation, the Web Engine serves the content based on the following:

- **Freshness of content**—The freshness of prefetched content is governed by a time-to-live (TTL) value set for the content in the delivery service configuration. The TTL specifies the rate at which content freshness is checked. For cached content, which is content ingested by means of the dynamic ingest or hybrid ingest method, the freshness check is performed by the Web Engine in compliance with RFC 2616. For expired cached content, the local copy is deleted and fresh content is ingested.
- **Rate of data transfer**—The rate at which the content is sent can be configured on a per-delivery basis. By default, LAN bandwidth is used.
- **Content completeness**—Prefetched content is stored locally in the CDS in its entirety. For cached content, there are two cases when the content is not complete:
  - The Web Engine process halts or the Service Engine experiences a failure in the process of caching the content. In this case, the subsequent request starts the cache fill anew.
  - The content is in the process of being cached by another request. In this case, the subsequent request is served from the cached content.

Authentication

The Web Engine supports a pass-through mode of authentication, whereby the origin server negotiates authentication and the Web Engine passes the requests and responses between the client device and the origin server. Content that requires authentication is not cached by the Service Engine, so all requests for authenticated content are retrieved from the origin server.

Service Rules

Service rules can be configured that dictate how the Web Engine responds when client requests match specific patterns. The patterns can be a domain or host name, certain header information, the request source IP address, or a Uniform Resource Identifier (URI). Some of the possible responding actions are to allow or block the request, generate or validate the URL signature, send an ICAP request to the policy server, or rewrite or redirect the URL.

Windows Media Engine

The Windows Media Engine uses Windows Media Technology (WMT), a set of streaming solutions for creating, distributing, and playing back digital media files across the Internet. WMT includes the following applications:

- Windows Media Player—End-user application
- Windows Media Server—Server and distribution application
- Windows Media Encoder—Encodes media files for distribution
- Windows Media Codec—Compression algorithm applied to live and on-demand content
- Windows Media Rights Manager (WMRM)—Encrypts content and manages user privileges

The Windows Media Engine streams Windows Media content, with the capability of acting both as a server and as a proxy. It streams prefetched content to the Windows Media Player, acts as a proxy for client requests, splits a live stream into multiple live streams, and caches content requested from remote servers.

Windows Media Engine acts as Windows Media Server for prefetched or cached content stored locally. The request is served by means of RTSP and HTTP. Windows Media Engine checks with the storage function on the Service Engine to see whether the content is stored locally; if the content is not found, the Windows Media Engine engages the Windows Media Proxy.

The WMT Proxy works like the cache fill operation in the Web Engine. There are two options:

- Hierarchical Caching Proxy—If content is not found locally, the Windows Media Engine checks the upstream Service Engines first before pulling the content from the origin server.
- Static Caching Proxy—The administrator statically configures Service Engines as upstream proxies.

The WMT Proxy accepts and serves streaming requests over RTSP and HTTP.

Fast Start

Fast Start provides data directly to the Windows Media Player buffer at speeds higher than the bit rate of the requested content. After the buffer is filled, prefetched, cached, or live content stream at the bit rate defined by the content stream format. Fast Start does not apply to content that is dynamically ingested. Only Windows Media 9 Players that connect to unicast streams using MMS-over-HTTP or RTSP can use Fast Start. The Fast Start feature is used only by clients that connect to a unicast stream. With live content, the Windows Media Engine needs to hold the content in its buffer for a few seconds. This buffer is used to serve Fast Start packets to subsequent clients that request the same stream as the initiating first client request. The first client triggers the process, with the subsequent clients benefitting from Fast Start.

Fast Cache

Fast Cache allows clients to buffer a much larger portion of the content before rendering it. Fast Cache is supported only for TCP. It is not supported for RTP. The Windows Media Engine streams content at a much higher data rate than specified by the stream format. For example, using Fast Cache, the Windows Media Engine can transmit a 128-kilobit per second (Kbps) stream at 700 Kbps. This allows the client to handle variable network conditions without perceptible impact on playback quality. Only MMS-over-HTTP and RTSP requests for prefetched or cached content support Fast Cache. The speed is determined by the client's maximum rate and the configured Fast Cache rate—whichever is smaller.

Fast Stream Start

The first client requesting a live stream often experiences the longest wait time for the content to begin playing. Users can experience long wait times because of the full RTSP or HTTP negotiation that is required to pull the live stream from the source. Delays can also occur if the edge Service Engine has not buffered enough stream data to fill the player's buffer at the time the content is requested. When the buffer is not filled, some data to the client might be sent at the linear stream rate, rather than at the Fast Start rate. With Fast Stream Start, when a live stream is primed, or scheduled and pulled, a live unicast-out stream is pulled from the origin server to a Service Engine before a client ever requests the stream. When the first request for the stream goes out, the stream is already in the delivery service.

Live Stream Splitting

Live stream splitting is a process whereby a single live stream from the origin server is split and shared across multiple streams, each serving a client that requested the stream. When the first client that requested the stream disconnects, the Windows Media Engine continues to serve the subsequent requesting clients until all requesting clients have disconnected. Live stream splitting using content that

is already stored locally is generally better than using content from the origin server; this is because the Service Engine is typically closer to the requesting clients, and therefore network bandwidth to the origin server is freed up.

Note When using Windows Media Server 2008 as the origin server, the source content type must be a playlist or encoder type.

Live stream splitting can either be unicast or multicast, depending on the configuration, capabilities and limitations of the network. The Windows Media Engine can receive and deliver Windows Media content over IP multicast or unicast transmission in the following combinations:

- Unicast-In Multicast-Out
- Multicast-In Multicast-Out
- Unicast-In Unicast-Out
- Multicast-In Unicast-Out

Multicast-Out

The Windows Media Engine can act as a multicast station to deliver multicast streams to client devices. The source of the stream can be multicast, unicast, or a local file. The station can be scheduled, continuous, or play once. The content can be either live or rebroadcast. The Windows Media Engine creates a Windows Media Station file (.nsc) that contains session information including the multicast IP address, port, time-to-live (TTL), and so on. The client requests the .nsc file using HTTP. Once the file is downloaded, the client parses it and sends an Internet Group Management Protocol (IGMP) join to receive the multicast stream. A client can start and stop the stream, but cannot pause, fast-forward, or rewind it. Multicast logging is a configurable option when setting up a multicast station. When multicast logging is enabled, the Windows Media Player automatically collects and sends the statistics to the multicast logging URL using the HTTP POST request method. The statistics collection is accomplished either by a remote server or by the Service Engine itself.

Unicast-Out

The Windows Media Engine can act as a broadcast publishing point to deliver live streams, prefetched/cached content, or content from dynamic ingest, to a requesting client. The source of the stream can be multicast, unicast, or a local file. The Windows Media Engine can also perform live stream splitting if more than one client requests the same content. The broadcast alias can be used to simulate an experience similar to viewing a TV program even if the source of the stream is a Video On Demand (VOD) file. A client can start and stop the stream but cannot pause, fast-forward, or rewind it. When a broadcast alias is configured, a client makes a request to the Windows Media Engine, which is acting as the Windows Media Server, and the Windows Media Engine checks to see whether the incoming stream is present. If it is, the Windows Media Engine joins the stream and splits it to the new client. If the request is the first client request for this stream, the Windows Media Engine sends the request to the origin server and then serves it to the new client.

Authentication

The Windows Media Engine supports pass-through authentication. The following authentication mechanisms are supported in pass-through mode:

- Anonymous
- NTLM
- Negotiate (Kerberos)
- Digest access authentication

With pass-through authentication, the Windows Media Engine establishes a tunnel between the client and the origin server so that the origin server can authenticate the client.

Bandwidth Management

Bandwidth management of Windows Media content can be controlled by setting limits for incoming and outgoing bandwidth and session bit rate and Fast Start maximum bandwidth. In addition, in the case of live streaming, contributing origin servers can by identified to allow incoming content to exceed the bandwidth check to support high demand scenarios. The Windows Media bandwidth management capabilities are described in Table 1-3.

Table 1-3 Bandwidth Management Capabilities

Bandwidth Management	Description
Incoming Bandwidth	The bandwidth for Windows Media content coming into the Service Engine, from either an upstream Service Engine or from the origin server.
Outgoing Bandwidth	The bandwidth for streaming Windows Media content to the end user from the Service Engine.
Incoming Session Bit Rate	The maximum bit rate per session that can be delivered to the Service Engine from the origin server or upstream Service Engine.
Outgoing Session Bit Rate	The maximum bit rate per session that can be delivered to a client.
Incoming Bandwidth Bypass List	The list of identified hosts allowed to bypass the incoming bandwidth check for broadcast or multicast live content.
Fast Start Maximum Bandwidth	Maximum bandwidth allowed per player when Fast Start is used to serve packets to each player. Increased bandwidth initially used by the Fast Start feature can overburden a network if many players connect to the stream at the same time. To reduce the risk of network congestion caused by the Fast Start feature, limit the amount of bandwidth the Fast Start feature uses to stream to each player.

Policy Server Integration

The Windows Media Engine uses HTTP and RTSP to send start, stop, and pause messages to the policy server.

Movie Streamer Engine

Note In Release 2.0, Movie Streamer Engine was in a demonstration state for live streaming, prefetched, cached, and dynamically cached content.
In Releases 2.1 on, live streaming is in full production. Prefetched, cached, and dynamically cached content remain in a demonstration state.
For details of live streaming performance for Movie Streamer Engine, please refer to the Release 2.1 performance bulletin.

The Movie Streamer Engine is an open-source, standards-based, streaming server that delivers hinted MPEG-4, hinted 3GP, and hinted MOV files to clients over the Internet and mobile networks using the industry-standard RTP and RTSP. Hinted files contain hint tracks, which store packetization information that tell the streaming server how to package content for streaming.

The Movie Streamer Engine is an RTSP streaming engine that supports Third Generation Partnership Project (3GPP) streaming files (.3gp). Support of 3GPP provides for the rich multimedia content over broadband mobile networks to multimedia-enabled cellular phones.

Note The streaming capability of Movie Streamer Engine only depends on the movie file format or stream transport type. It is independent of codec types. Movie Streamer supports any client player that can fetch media streams by way of RTSP or RTP. However, the client player must have the correct codec in order to render the stream correctly.

The Movie Streamer Engine can act as both a server and a proxy. It streams prefetched or RTSP-cached content to RTSP clients, acts as a proxy for client requests, splits a live stream into multiple live streams, and caches content requested from remote servers.

After the RTSP request comes into the Movie Streamer, the URI in the RTSP request is modified to reflect the result of the mobile capability exchange. The Movie Streamer checks with the storage function on the Service Engine to see whether the content is stored locally. If the content is not found or if an RTSP-cached content version needs freshness validation, the Movie Streamer engages the Movie Streamer proxy.

In the case of an RTSP-cached content version verification, the Movie Streamer proxy forwards the DESCRIBE request to the origin server for a response containing the Last-Modified-Time header in the response. If the Last-Modified-Time matches the cached version, the Movie Streamer streams the cached content; otherwise, the Movie Streamer proxy forwards the request to the origin server for RTSP negotiation. Then, a client session and a server session are created.

- The server session is responsible for connecting to the origin server to fetch the content and cache it locally. The server session generates the media cache file and the linear hint files.
- The client session is responsible for streaming the locally cached file to the client.
- The client and server sessions are separated so that multiple server sessions can be spawned for the same URL to cache content from different starting points or at faster speeds, or both. This increases the speed of fetching the content. The client session starts to stream from the cached content that the server session is writing.

The Movie Streamer proxy works like the cache fill operation in the Web Engine and the Windows Media Engine. There are two options:

1. Hierarchical Caching Proxy—If content is not found locally, the Movie Streamer checks the upstream Service Engines first before pulling the content from origin server.
2. Static Caching Proxy—The administrator statically configures Service Engines as upstream proxies.

The Movie Streamer supports basic pass-through proxy mode for certain conditions where caching cannot be performed. Such conditions include, but are not limited to, the Service Engine running out of disk space.

Transport Types

Prefetched content can be delivered by the non-accelerated method or the accelerated method. Non-prefetched content (proxied or cached content) is always delivered by the accelerated method. The content is delivered to the client device by means of one of the following mechanisms:

- **Non-Accelerated**—This method has limited concurrent streams and total throughput, but supports many transport formats. The non-accelerated method supports the following transport formats:
  - RTP over UDP
  - Reliable UDP

- **Accelerated**—This method supports only RTP over UDP. Content must be reprocessed by the Movie Streamer Linear Hinter. The linear hinter process can be initiated manually by the administrator or dynamically triggered by the first request for the content.

The Movie Streamer Linear Hinter process may take a while, so the first request that triggers this process is served by means of the non-accelerated method. All subsequent requests are served by the accelerated method.

The first client request for content that requires proxying or caching experiences a delay, because all proxying and caching requires the accelerated method.

Live Stream

The Movie Streamer Engine supports multicast reference URLs (Announce URLs) for programs that are created through the Internet Streaming CDSM. The multicast reference URL, which is in the form of http://*Service Engine IP address*/*Program ID*.sdp, is resolved by the Movie Streamers that are serving the live program.

QuickTime live typically has a UDP socket pair (for RTP and RTCP) per track, and each client session typically has two tracks (audio and video).

Note The following rules apply to live splitting:

1. For unicast streaming, the client request must be sent by means of RTSP.
2. For multicast streaming, the client request must be sent by means of HTTP.

Authentication

The Movie Streamer Engine supports the following authentication modes:

- Basic
- Digest

URL Signing

Support of URL signing for Movie Streamer content requests is a Release 2.2 feature. For more information see the "URL Signing" subsection, under the "Flash Media Streaming Engine"section on page 1-17.

Flash Media Streaming Engine

Note Flash Media Streaming is a Release 2.1 and Release 2.2 feature.

The Flash Media Streaming Engine incorporates the Adobe Flash Media Server technology into the CDS platform. The Flash Media Streaming Engine is capable of hosting Flash Media Server applications, such as VOD (prefetched content, or dynamic or hybrid ingested content) and live streaming, that are developed using ActionScripts.

Upon receiving a client request for VOD content, the edge Service Engine does the following:

- If the content is present, the edge Service Engine streams it using RTMP.
- If the content is not present, the edge Service Engine uses HTTP to fetch the content from the origin server and serves it using RTMP.

No client information is sent to the origin server. No per-client control connection is present between edge to origin server for VOD streaming.

Note The Cisco CDS Flash Media Streaming Engine supports the Adobe Flash Media Rights Management Server (FMRMS) for VOD content; it is not supported for live streaming. Adobe FMRMS protects media content delivered to Adobe Media Player and Adobe AIR applications. FMRMS is also available for proxied content, if Adobe supports the content type. For more information about the Adobe Flash Media Rights Management Server, see www.adobe.com.

Note Release 2.3 supports the Adobe Flash Media Server Administration APIs and the Administration Console that was built using the Administration APIs. These APIs can be used to monitor and manage the Adobe Flash Media Server running on a Cisco CDS Service Engine. See the "Configuring Flash Media Streaming" section on page 4-42 for more information.

HTTP Requests

Flash Media Streaming encompasses all flash applications, from simple Flash Video (FLV) files to more complex Small Web Format (SWF) files. All HTTP client requests for SWF files, that are redirected to a Service Engine by the Service Router, are handled by the Web Engine. The Web Engine, using HTTP, serves the request from locally stored content in the CDS or from any upstream Service Engine or origin server. See the "Web Engine" section on page 1-9 for more information.

RTMP Requests

The SWF file is a compiled application that runs on the Adobe Flash Player, and may contain Real Time Media Protocol (RTMP) calls to FLV, MPEG-4 (H.264), or MP3 files. RTMP calls, in the form of URL requests, are routed to a Service Engine by the Service Router.

Flash Media Streaming supports RTMP and RTMPE on port 1935 only. RTMPE is the secure flash streaming technology from Adobe. Encrypted RTMP (RTMPE) is enabled on Flash Media Streaming by default, and allows you to send streams over an encrypted connection without requiring certificate management.

In Release 2.3, Flash Media Streaming also supports RTMPT and RTMPTE on port 80. RTMP Tunneled (RTMPT) encapsulates the RTMP data within HTTP requests in order to traverse firewalls. RTMP Tunneled Encrypted (RTMPTE) encrypts the communication channel, tunneling over HTTP.

Note The Service Router uses RTMP redirection to direct the client's Flash Player to the best Service Engine based on load balancing and resiliency. RTMP redirections are supported only by Adobe Flash Player 9. All older Flash Players do not support RTMP redirection.

Note For VOD streams, all RTMP calls in the SWF file must be in the following format:

```
rtmp://rfqdn/vod/path/foo.flv
```

In this format, *rfqdn* is the routing domain name of the Service Router, *vod* is the required directory, and *path* is the directory path to the content file that conforms to the standard URL specification.

For prefetched and cached content, the Flash Media Streaming Engine uses RTMP or RTMPE over port 1935. In Release 2.3, the Flash Media Streaming Engine also supports RTMPT and RTMPTE over port 80. For content that is not found locally, the Flash Media Streaming Engine communicates with the Web Engine, that in turn communicates with the upstream Service Engine for cache fill operations. See the

"Cache Fill Operations" section on page 1-9. This interaction uses HTTP. Once the content is in the process of being retrieved by the Web Engine, the Flash Media Streaming Engine uses RTMP to begin streaming the content.

The following describes the characteristics of caching content using HTTP for RTMP client requests;

1. Origin server-based cache validation is still honored for the cached content.
2. Client-side Web Engine rules are bypassed for the RTMP client request.
3. If HTTP headers from the origin server have the "no-cache" attribute set, content is not cached, and transparent proxy is performed to stream RTMP.
4. Transparent proxy from HTTP to RTMP is supported. Flash Media Streaming Engine begins RTMP streaming while content is still being fetched using HTTP proxy mode.

Any HTTP configuration that prevents content from being cached still applies for RTMP requests. The Flash Media Streaming Engine uses multiple HTTP-based range requests in such cases.

Flash Media Streaming Proxy

The Flash Media Streaming Engine can deliver content acting as an origin server or as a proxy server. The Flash Media Streaming Engine acts as a proxy server when content cannot be cached due to the origin server's configuration or due to the Service Engine's Web Engine configuration. Content is ingested and distributed using HTTP, whether the client request for the content used HTTP or RTMP.

Note Any content that does not contain "live" or "vod" in the path is automatically proxied.

Unicast Streaming

The Flash Media Streaming Engine supports unicast flash streaming.

URL Signing

Note Support of URL signing for Flash Media Streaming content requests is a Release 2.2 feature.

Flash Media Streaming supports signed URLs, which adds additional security. The URL signature generation is based on a key that is a shared secret between the component generating the URL signature and the component validating the URL signature. The URL signature can be generated by the Service Engine, another component external to the Service Engine, or the web portal.

For more information about the URL signatures, see the "Configuring URL Signing" section on page 4-23.

Codecs

The Flash Media Streaming codec that is supported in Release 2.1 is On2 VP6. The Flash Media Streaming codecs supported in Release 2.2, in addition to On2 VP6, are listed in Table 1-4.

Table 1-4 Codecs Supported in CDS Releases 2.2 and 2.3 for Flash Media Streaming

Standard	Details
ISO/IEC 14496-3	MPEG-4 Part 3, also known as AAC+, HE-AAC. A set of compression codecs for perpetual coding of audio signals, including some variations of Advanced Audio Coding (AAC), as well as AAC Main, AAC LC, and SBR.
ISO/IEC 14496-10	Advanced Video Coding (AVC), also known as H.264/AVC. All levels of applications are supported, Base (BP), Main (MP), High (HiP), High 10 (Hi10P), and High 4:2:2 Profile (Hi422P). This standard is technically identical to the ITU-T H.264 standard.
ISO/IEC 14496-12	ISO Base Media File Format. A file format for storing media content containing one audio track (either ISO/IEC 14496-3 [AACPlus] or MP3), and one video track (either ISO/IEC 14496-10 [H.264 or AVC] or VP6).
3GPP TS 26.245	Time text format

Live Streaming

Note Live streaming using Flash Media Streaming is a Release 2.2 feature.

Flash Media Streaming uses RTMP to stream live content by means of dynamic proxy. Configuration of live or rebroadcast programs is not required. When the first client requests live streaming content, the stream is created. There are no limits to the number of live streams other than the system load. Live streaming uses distributed content routing to distribute streams across multiple Service Engines.

Upon receiving a client request for live content, the edge Service Engine does the following:

- If the live stream is already present, the edge Service Engine attaches the new client to the existing stream. No message is sent to the origin server and no connection is set up.
- If the live stream is not present, CDS creates a connection to the origin server to get the stream. No client information is sent to the origin server.

No per-client control connection is present between edge to origin server for live streaming.

For Flash Media Streaming, a delivery service can be used for prefetched content, cached content, dynamically cached content, and live content. Because Flash Media Streaming uses dynamic proxy to stream live content, no disk space is used to store content. A Service Engine can act as the origin server origin server for streaming live content, provided the SE designated as the origin server is not assigned to the delivery service that is streaming the live content.

The Flash Media Streaming Engine automatically retries a connection to the origin server if the upstream live-splitting connection fails. This failover does not require any additional retries from the client side. Clients see a subsecond buffering, after which video continues to play. This feature does not address failover when the Service Engine that is streaming to the client fails. The primary advantage is increased resiliency in the CDS infrastructure. In other words, if a Service Engine fails, the downstream Service Engine automatically connects to the origin server.

The Adobe Flash Media Encoder can publish the streams to any Adobe Flash Media Server acting as the origin server. Clients use the RFQDN to get the live content. The request from the client for "streamname" is mapped to origin_appinst_streamname internally in the CDS to differentiate between two streams with the same name in two different delivery services.

Note All RTMP calls for live content in the SWF file must be in the following format:
`rtmp://rfqdn/live/path/foo.flv`
In this format, *rfqdn* is the routing domain name of the Service Router, *live* is the required directory, and *path* is the directory path to the content file that conforms to the standard URL specification.

Flash Media Streaming supports live stream splitting. For more information about live stream splitting, see the "Live Stream Splitting" section on page 1-11.

Interactive Applications

Release 2.3 supports pass-through (proxy) support for interactive applications (non-VOD or non-live). The interactive applications are hosted on a Flash Media Interactive Server that is external to the CDS.

Direct routing from the Service Engine, acting as the Flash Media Streaming edge server proxy, to the origin server (the Flash Media Interactive Server) is supported. Using the delivery service framework, the origin server is abstracted from the client request by using the Service Router Domain Name (SRDN), which resolves to the Service Engine that accepts the user connection and forwards the request to the origin server. In Release 2.3, Flash Media Streaming includes the edge server (proxy) mode, and by default, all non-live and non-VOD applications are proxied by using the edge server. Flash Media Streaming selectively picks connections for processing in edge server mode and aggregates connections to the origin servers.

CDS supports implicit URI as the method that allows the client to connect with the edge server without exposing the origin server. The URI would look like this: `rtmp://edge1.fms.com/ondemand`.

Request routing based on SWF files or using RTMP redirection is supported. However, RTMP redirection requires more changes in the client code. SWF file-based redirection is recommended. SWF redirection works as follows:

1. SWF files and associated HTML pages are either prefetched or hosted in the origin server.
2. Client uses a web browser to access the HTML page, which also loads the SWF file.
3. The SWF file is accessed using the SRDN.
4. The Service Router redirects the request to a Service Engine.
5. The SWF file is downloaded to the web browser.
6. The ActionScript in the SWF file attempts to connect to the same host from where the SWF file was downloaded. This is an RTMP connection that reaches the Service Engine.
7. The Service Engine checks for the application type in the URI, if it is not VOD or live, the processing is moved to the edge server mode and the connection is forwarded to the origin server.
8. The Service Engine tunnels the data between the client and the origin server.

Note Changes to a delivery service do not affect existing connections to the Flash Media Interactive Server (origin server). Only new connections are affected by changes to a delivery service.

Service Router

The Service Router mediates requests from the client devices and redirects the requests to the most appropriate Service Engine. It monitors the load of the devices and does automatic load balancing.

The Service Router is the authoritative Domain Name System (DNS) server for the routed request for the Fully Qualified Domain Name (FQDN) of the origin server. In other words, the Service Router responds to any DNS queries for that domain. The Service Router chooses the Service Engine based on two scenarios:

- Client is directly connected to the service provider's network (on-net).
- Client is roaming outside the home network (off-net).

When clients are connected to the service provider's network, the Service Engine is chosen based on the requested FQDN, the client's IP address, and the responsiveness of the Service Engine. The Service Router compares the source IP address of the client against a table of address ranges assigned to the Service Engines, known as the *Coverage Zone file*. The Coverage Zone file provides information on the proximity of the client to the Service Engine based on each client's IP address.

If the client is not connected to the service provider's network and location-based routing is enabled, the Service Engine that is geographically closest to the client is selected. The Service Router uses a Geo-Location server to find the geographical location of the client. Geographical distance is calculated between the client and the Service Engine. The closest Service Engine is chosen based on the distance.

IP-Based Redirection

Note IP-based redirection is a Release 2.3 feature.

In Release 2.3, there are two ways for client requests to get routed to the Service Router and on to the Service Engine:

- Router Fully Qualified Domain Name (RFQDN) redirection
- IP-based redirection

RFQDN redirection is the default configuration. With RFQDN redirection, client requests are resolved to the Service Router by the DNS server and the Service Router redirects the request to the Service Engine based on route tables created from the Coverage Zone file and the current load of the Service Engines. The redirected URL is http://SENAME.SE.RFQDN/relative_path_of_content, where SENAME is the hostname of the Service Engine.

Note The redirected URL for Flash Media Streaming requests is rtmp://SENAME.SE.RFQDN/application_name/encoded (relative_path_of_streamname), where SENAME is the hostname of the Service Engine.

When IP-based redirection is enabled, the Service Router uses the IP address of the Service Engine in the URL instead of the hostname. The redirected URL is http://<se ip addr>/ipfwd/<rfqdn>/<path>. The IP-based redirection method avoids the extra DNS lookup that was required in the RFQDN redirection.

Service Router Workflow

The Service Router workflow for clients connected to the service provider's network is as follows:

1. The client sends the DNS query for the routed FQDN to the local DNS server.
2. The DNS server replies with the Service Router IP address.
3. The client issues an HTTP or RTSP request to the Service Router.
4. If the Service Router finds the client's subnet in the Coverage Zone file, the following occurs:
   a. The Service Router chooses the appropriate Service Engine and performs a protocol-specific redirection.
   b. The client issues an HTTP or RTSP request to the Service Engine.
   c. The Service Engine serves the content.

   If the Service Router does not find the client's subnet in the Coverage Zone file and location-based routing has been enabled, the following occurs:
   a. The Service Router communicates with a Geo-Location server and gets the geographical coordinates of the client's IP address.
   b. The distance is calculated between the client and the Service Engines, and the Service Engine closest to the client is selected.
   c. The Service Router performs a protocol-specific redirection with the closest Service Engine.
   d. The client issues an HTTP or RTSP request to the Service Engine.
   e. The Service Engine serves the content.

When a Service Router is registered with the Internet Streaming CDSM, the CDSM propagates the Service Router's IP address to all the registered devices. The Service Engine sends a keep-alive message to the Service Router on a periodic interval, which consists of information about the disk health, whether the Internet Streamer application is enabled or not, and the load of the Internet Streamer application on the Service Engine. The Service Router uses the Service Engine's load and liveness information for generating the routes.

The Internet Streamer CDS can have more than one Service Router in order to support Service Router failover. In line with failover, the DNS server should be configured with multiple Service Routers for the same routed FQDN.

Note DNS entries for all FQDNs must be delegated to the Service Router. In the DNS server's database file, a name server record must be entered for each FQDN that routes to the Service Router.

Coverage Zone File

When a Service Engine is registered to the Internet Streaming CDSM, it is assigned a default Coverage Zone file that is created by the CDSM using the interface IP address of the Service Engine. The default Coverage Zone file can be unassigned, and a custom coverage zone can be created using the Coverage Zone file.

A Coverage Zone file is an XML file containing coverage zone entries for each client IP address range, the Service Engine serving that range, the latitude and longitude of the Service Engine, and a metric value. The Coverage Zone file can be referenced by a URL and imported into the Internet Streaming CDSM, or uploaded from a local machine. The Coverage Zone file can be set as the default for a specific Service Router or for all Service Routers in the CDS network.

When content is requested by a client, the Service Router checks the client's IP address to find the coverage zone that contains that IP address. The Service Router then selects the Service Engine that serves this coverage zone. If a specific IP address is in multiple coverage zones, the one with the more specific range is selected. If no match is found in the coverage zone data and if location-based routing is enabled on the Service Router, the Service Router looks up the best Service Engine closest to the client. If the Service Router is unable to redirect the request, the Service Router sends an error response to the client.

A coverage zone can be associated with one or more Service Engines. Each Service Engine can have its own unique coverage zone, or the Service Engines can be associated with more than one coverage zone and have over lapping coverage zones. In Figure 1-2, all Service Engines serve Coverage Zone 1, and Service Engine 1 is specifically associated with Coverage Zone 2, a subset of Coverage Zone 1.

Figure 1-2 Coverage Zone Example

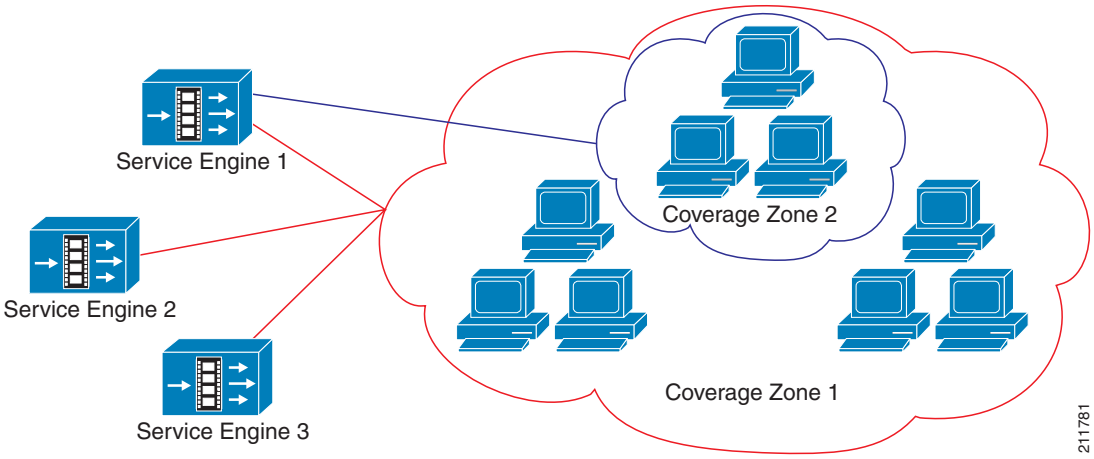

If a coverage zone is served by multiple Service Engines, all Service Engines are put in the routing table. The metric value, entered in the Coverage Zone file, indicates the proximity of the Service Engine to the client. When multiple Service Engines serving a coverage zone are on the same subnet and have the same metric value, and load-based routing is not enabled, the Service Router uses round-robin routing to redirect the client. If load-based routing is enabled, the load of the Service Engines are used to determine the best Service Engine to redirect the client.

Routing Methods

The Service Router chooses the best Service Engine based on whether the Service Engine is participating in the delivery service for which the origin server matches that of the requested domain and whether the Service Engine is assigned to serve the client's network region.

The Service Router uses four methods:

- Simplified hybrid routing (round-robin-based routing) (Releases 2.0 and 2.1)
- Load-based routing (least loaded)
- Last-resort routing (all Service Engines are overloaded)
- Location-based routing (off-net clients)
- Service aware routing
- Content-based routing

Note In Release 2.2, the keepalive messages between the Service Router and Service Engine are transmitted and received on port 2323. However, the Release 2.2 software interoperates with older software releases that do not use port 2323 for keepalive messages. If a firewall is configured between the Service Engine and the Service Router, port 2323 (UDP) must be opened for the keepalive message to go through.

Simplified Hybrid Routing

Note This feature is supported only in Releases 2.0 and 2.1.

The Service Router uses simplified hybrid routing to route requests to Service Engines. In simplified hybrid routing, the Service Engines are chosen in a round-robin fashion if multiple Service Engines are serving the client network region.

Load-Based Routing

Note This is a Release 2.0 and Release 2.1 feature. In Release 2.2 and later, load-based routing is enabled by default and cannot be disabled.

If load-based routing is enabled, the routing decision is made according to the capacity and load of the Service Engines instead of selecting the Service Engines in a round-robin fashion. This routing method is enabled by default.

The load of the Service Engine is determined by different parameters, such as processor use, memory usage, disk usage, the number of current Windows Media streams being served, and so on. The current load is compared with the thresholds configured for the Service Engine. If a threshold has been exceeded for a Service Engine it is excluded from the routing table.

Last-Resort Routing

Last-resort routing is applicable when load-based routing is enabled and all Service Engines have exceeded their thresholds or all Service Engines in the domain are offline. The Service Router can redirect requests to a configurable alternate domain when all Service Engines serving a client network region are overloaded.

Last-resort routing works dynamically when Service Engines are overloaded or deactivated. When the load of one or more Service Engines in the original host domain is reduced below threshold limits or the Service Engines are reactivated, new requests are routed to the original host domain automatically.

If a last-resort domain is not configured and the Service Engine thresholds are exceeded, requests are redirected to the origin server.

Last-resort routing supports requests from RTSP, HTTP (including MMS-over-HTTP), and RTMP clients.

Location-Based Routing

Location-based routing is used for off-net clients. Off-net clients are clients that are not directly connected to the service provider's network. Location-based routing is designed to work with load-based routing. When both are enabled, the Service Router first looks up the client's IP address in the Coverage Zone file. Then, if there is no subnet match, the client's geographical location is compared to the

geographical location of the Service Engines listed in the Coverage Zone file, and the closest and least-loaded Service Engine is selected. Geographically locating a client is used when users roam outside of their home network.

In order to provide routing to off-net clients, the Service Router communicates with a Geo-Location server, which maps IP addresses to a geographic location. For redundancy, the Internet Streaming CDSM can be configured with a primary and secondary Geo-Location server.

The Geo-Location Server identifies the geographical location of an off-net client by the latitude and longitude of the client. The Service Router compares the client's location with the location of the Service Engines participating in that delivery service and chooses the best Service Engine to serve the content.

Service Aware Routing

Note This is a Release 2.2 feature. Service aware routing is always enabled and is not configurable.

In service aware routing, the Service Router redirects the request to the Service Engine that has the required protocol engine enabled, the required protocol engine is functioning properly and has not exceeded its threshold, and the SE has not exceeded its thresholds as configured. See the "Setting Service Monitor Thresholds" section on page 4-77 for more information.

When a request reaches the Service Router, the Service Router generates a hash from the URI. The Service Router first generates a list of Service Engines to best serve the request based on service aware routing. The Service Router then reorders the list based on the hash and selects the best Service Engine. Because the hash generated for the same URI is equal, typically the same Service Engine is selected. If the Service Engine is overloaded, the next Service Engine in the list is selected.

Note For service aware routing, some of the services running on a Service Engine are protocol based. When protocol-based services associated with a protocol engine are stopped on a Service Engine, the Service Router excludes this Service Engine from the list of possible Service Engines that can serve requests for this type of content. The Service Router identifies the protocol engine that will serve the request based on the user-agent in the request. For example, if some Windows Media Engine-related services are stopped, the Service Engine can still serve Web Engine requests. However, if the request for Web Engine content is sent from a Windows Media Player, the Service Router excludes the Service Engine from the list of possible Service Engines that can serve the request.

Note For service aware routing, when a threshold is exceeded for all Service Engines, the Service Engine redirects the client request to the origin server if an alternate domain is not configured. If an alternate domain is configured, the alternate domain takes precedence. For a managed-live URL, if the origin server does not match the source of the live program, the above case fails. For the above case to work, the origin server host must be configured to match the live program source. In addition, the origin server stream name must be the same as the live program name.

Content-Based Routing

Note This is a Release 2.2 feature.

In content-based routing, the Service Router redirects the request based on the Uniform Resource Identifier (URI). Requests for the same URI are redirected to the same Service Engine, provided the Service Engine's thresholds are not exceeded.

The number of redundant copies of content is configurable for content-based routing. The same content can be stored in more than one Service Engine if the number of redundant copies is set to more than one. Redundancy is used to maximize the cache hit ratio. If redundancy is configured with more than one copy, multiple Service Engines are picked for a request with the same URI hash.

Content-based routing is best suited for cache, prefetched, and live program requests in order to maximize the cache hit ratio.

Request Redirection

The Service Router supports the following redirections:

- HTTP ASX Redirection — Used if the requested file has an.asx extension. This redirection method is used for Windows Media Technology.
- HTTP 302 Redirection — Used if the protocol is HTTP and the file extension is not .asx. This is the native HTTP redirection.
- RTSP 302 Redirection — Used if the protocol is RTSP and the client is QuickTime or Windows Media. This is the native RTSP redirection.
- RTMP 302 Redirection — Used if the protocol is RTMP and the client is Adobe Flash Player, Adobe Media Player, or Adobe Flash Lite Player.

Content Delivery System Manager

The Internet Streaming Content Delivery System Manager (CDSM) is a web browser-based user interface. The Internet Streaming CDSM allows the administrator to configure, manage, and monitor delivery services and devices in the Cisco Content Delivery System (CDS). Application programming interfaces (APIs) are provided for backoffice integration with the Internet Streaming CDSM.

Authentication, Authorization, and Accounting

The Internet Streaming CDSM uses HTTPS to secure the administrator's session. Multiple users can perform administrative operations by using the Internet Streaming CDSM. The administrator can configure certain users to have either view-only rights for monitoring the CDS, or full rights that allow configuration changes as well as monitoring capabilities.

User accounts and groups can be added to the Internet Streaming CDSM and given roles and rights for accessing configuration information. It is also possible to segregate and group objects and give access to a limited group of users.

User authentication can be configured to use RADIUS servers when available, otherwise the Internet Streaming CDSM provides its own authentication server.

The CDS-wide policy and status information is maintained in a relational database on the Internet Streaming CDSM. This information is propagated and synchronized with all devices in the CDS network.

As part of the network management process, the administrator can perform basic administration operations on the Internet Streaming CDSM database, including backup and restore.

Device Management

The Internet Streaming CDSM sends device configuration changes to the selected device or group of devices once the change has been submitted. The device sends any configuration changes that were made locally to the CDSM, and also provides periodic status information.

Devices can be organized into user-defined device groups, which allow administrators to apply configuration changes and perform other group operations on multiple devices simultaneously. Because a device can belong to multiple device groups, this reduces the management overhead of the administrator. Device groups allow for a single instance of management thus eliminating the need to repeat the same step for each device.

The Internet Streaming CDSM also provides an automated workflow to apply software upgrades to the devices in a device group.

Higher Storage Utilization of CDS

Storage across multiple Service Engines is virtually divided into buckets where each Service Engine serves only a subset of the total content. Both the local storage and RAM of the Service Engines can function as an aggregated distributed service, providing unlimited scalability. Linear scaling of the CDS storage is accomplished by adding more Service Engines to one location. This addresses the demands of the "Long Tail" use case relevant to the Service Engines. The Long Tail is the realization that the sum of many small markets is worth as much, if not more, than a few large markets. Long-tail distribution is the possibility that extremely infrequent occurrences in traffic are more likely than anticipated.

This higher storage utilization provides the following:

- Overall better system performance
- Higher in-memory cache hit ratio
- Deterministic resiliency in case of failures or overload due to very popular content (This is useful when customers have live, prefetched, and cached assets more than 4.5 terabytes of content on one Service Engine.)

The content distribution is resilient and stateless. If the load of all content mapped to one Service Engine increases, the load is automatically spread to other Service Engines without requiring any administrator intervention.

Delivery Services Management

The Internet Streaming CDSM provides the configuration and monitoring of delivery services, which defines how content is ingested, stored, cached, and published. The Internet Streaming CDSM provides the Service Engines with information about the delivery services and which Service Engines are participating in the delivery service.

In addition to using the Internet Streaming CDSM to define delivery services, an XML file called a *Manifest file* can be used to define a delivery service. The Manifest file and APIs serve as the basis for backoffice integration. For more information about the Manifest file, see the "Manifest File" section on page 2-5.

Resiliency and Redundancy

A CDS that is designed with full redundancy and no single point of failure includes redundant Internet Streaming CDSMs and Service Routers. The redundancy mechanisms for the Content Acquirer and Internet Streamer applications running on the Service Engines operate differently.

Content Acquirer Redundancy

In the event of a primary failure on the Content Acquirer, the failover mechanism supports the election of a backup Content Acquirer. A failover requires that both the primary and backup Content Acquirer be located in the root location of the delivery service.

Live Programs

If the Content Acquirer receives a live program as a multicast stream from the origin server, upon failure of the primary, the backup Content Acquirer assumes control of that program's streaming and the program continues without interruption. This process is transparent to the end user. When the primary Content Acquirer comes back online, it receives the live stream from the active secondary Content Acquirer and does not fall back (regain its primary status) until the live program has finished or has been restarted.

If the Content Acquirer receives the program as a unicast stream from the origin server, the failover mechanism is not supported. If the primary Content Acquirer fails while a program is playing, the person viewing the program must re-request the program.

Internet Streamer Redundancy

If a Service Engine running the Internet Streamer application fails, the Service Router stops receiving keep-alive messages from that Service Engine. When a new request comes in, the Service Router does not redirect the request to that Service Engine; instead, it redirects the request to other Service Engines within the same delivery service. All the existing sessions on the failed Service Engine terminate and the affected end users must re-request the content.

Service Router Redundancy

If the CDS network is designed with multiple Service Routers, all Service Routers are aware of all Service Engines in the CDS. The DNS servers must be configured with multiple Service Routers and the failover is handled by the DNS servers.

Internet Streaming CDSM Redundancy

The Internet Streaming CDSM can operate in two different roles: primary and standby. The primary role is the default. There can only be one primary active in the CDS network; however, you can have any number of Internet Streaming CDSMs operating in standby to provide redundancy and failover capability.

Primary and standby CDSMs must be running the same version of software. We recommend that the standby CDSM be upgraded first, followed by the primary CDSM.

The Internet Streaming CDSM design principle is that the management device is never in the service delivery path. When the CDSM fails, the rest of the CDS continues to operate. A CDSM failure does not affect any services delivered to end users, and all content ingest continues. The only negative effect is

that the administrator cannot change configurations or monitor the CDS. As soon as a failure to connect to the CDSM is noticed, the administrator can activate the standby CDSM. For information on making the standby CDSM the primary CDSM, see the "Changing a Standby to a Primary CDSM" section on page 3-9.

CHAPTER 2

Network Design

Provisioning the Cisco CDS consists of two stages:

- Register the devices to the Internet Streaming CDSM and define the network topology and device groups.
- Configure the delivery services that deliver content to the clients.

This chapter describes the details of the two stages of provisioning a Cisco CDS network and how metadata and content flow through the Cisco CDS. This chapter has the following major topics:

- Cisco CDS Topology, page 2-1
- Delivery Service, page 2-3
- Service Workflow, page 2-7
- Programs, page 2-9

Note In order to achieve the best throughput, we recommend you configure a port channel for the four Gigabit Ethernet ports on the line card. For more information, see the "Configuring Port Channel" section on page E-1.

Cisco CDS Topology

In the Cisco CDS the Service Engines are grouped together into locations, such that a Location Tree is a set of locations organized in the form of a tree. The Location Tree represents the network topology configuration that is based on parent-child relationships. Locations are well connected and have similar connectivity properties to the outside world. A location generally implies topological proximity. Each location can have a parent relationship and multiple child relationships, such that each location can have zero to one parent locations and zero to many child locations. These relationships guide how content flows among locations but does not restrict content flow in any direction.

Locations are also classified into tiers. Each tier consists of locations belonging to the same tier. All locations with no parents belong to Tier 1. All locations that are children of Tier 1 locations belong to Tier 2.

The Cisco CDS can consist of one or more topological Location Trees. A Cisco CDS network is limited by the maximum depth of four tiers.

Figure 2-1 illustrates two location trees, with the parent-child relationship of each location indicated by a solid line and each tier indicated by a dotted line.

Figure 2-1 Location Trees Example

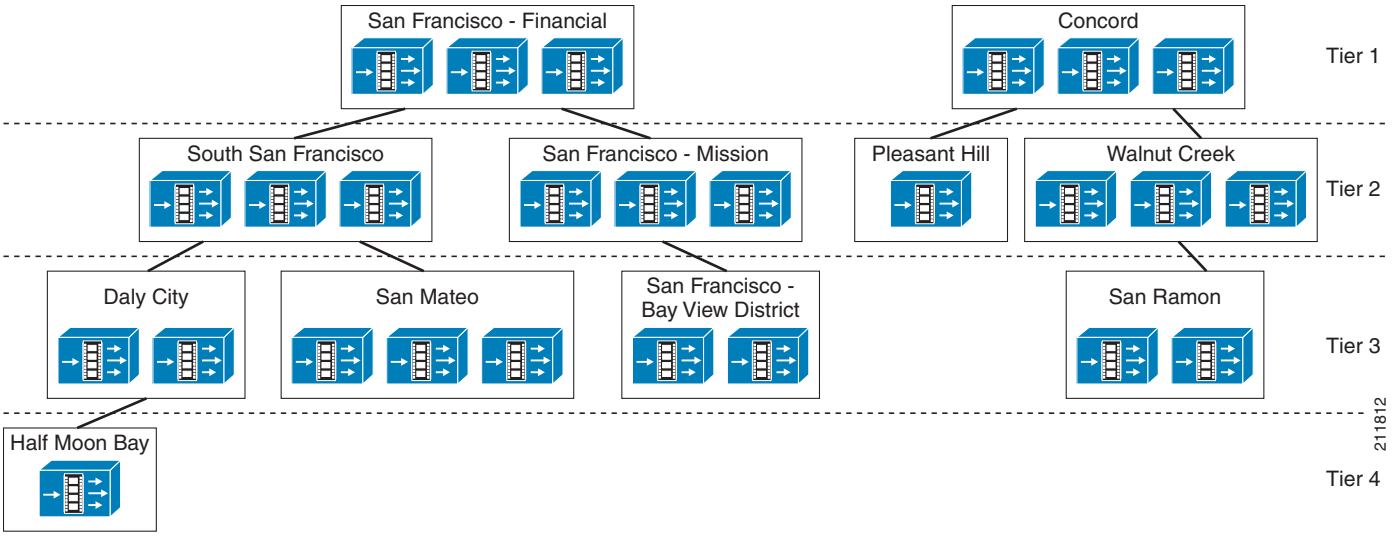

The Location Trees define preferred distribution routes. The Tier 1 locations are located closest to the Internet or backbone. Tier 1 locations can communicate with all other Tier 1 locations.

Note The CDS does not support network address translation (NAT) configuration, where one or more CDEs are behind the NAT device or firewall. The workaround for this, if your CDS network is behind a firewall, is to configure each internal and external IP address pair with the same IP address.

The CDS does support clients that are behind a NAT device or firewall that have shared external IP addresses. In other words, there could be a firewall between the CDS network and the client device. However, the NAT device or firewall must support RTP/RTSP.

Device Groups

Device groups offer a way to group similar devices and configure all the devices in a group at one time. Service Engines can be assigned to multiple device groups when the Device Group Overlap feature is enabled.

A device in a device group can have individual settings different from other devices in the group, and its settings can revert back to the group settings. The last configuration submitted for the device, whether group or individual, is the configuration the device uses.

In addition to group configuration and assignment, the CDSM allows the following:

- Hiding configuration pages of a device group
- Adding all newly activated devices to a device group
- Forcing device group settings onto all devices assigned to a group

A device can be assigned to a device group in one of two ways:

1. From the Device Assignment page
2. From the Device Group Assignment page

Baseline Groups

A baseline group is a special type of device group that denotes a group of devices for a particular service. There are three baseline groups:

- Web Baseline Group—Used for web-based content
- Video Baseline Group—Used for video content
- Platform Baseline Group—Used for platform-specific configurations

A device group can be configured as a baseline group. A device can be assigned to a baseline group in the following three ways:

1. From the Device home page.
2. From the Device Assignment page.
3. From the Device Group Assignment page.

Delivery Service

A delivery service is a configuration that defines how content is acquired, distributed, and stored in advance of a client request (prefetch), and after a client request (cached). Content from a single origin server is mapped to a set of devices by means of a delivery service. Content objects associated with a specific delivery service have a common domain name; in other words, the content in a specified delivery service resides in a single location on an origin server. Each delivery service maps service routing domain names to origin servers one-to-one for Service Router DNS interception.

For each delivery service, there is only one Content Acquirer but multiple Service Engines. The location that has the Content Acquirer for a delivery service is called the *root location*. Other Service Engines in the root location that are assigned to the same delivery service can act as backup Content Acquirers if the configured Content Acquirer fails.

Delivery services form logical routes for content to travel from an origin server through the Content Acquirer to all the Service Engines in the delivery service. Logical routes for content distribution are based on the device location hierarchy or Location Tree.

The content distribution route follows the general tree structure of the Location Tree, where content is distributed from the root of the tree (Content Acquirer) to the branches (Service Engines associated with the delivery service). A delivery service distribution tree is constructed for each delivery service.

By excluding it from the Coverage Zone file, a Service Engine in a delivery service can be configured only to forward content and metadata, and not deliver the content to client devices.

Figure 2-2 shows an example of a delivery service distribution tree. The Service Engines participating in the delivery service are marked in red. Possible content and metadata routes are indicated by red lines. The actual route may differ among the participating Service Engines as determined by the Service Router routing method.

Figure 2-2 Delivery Service Distribution Tree Example

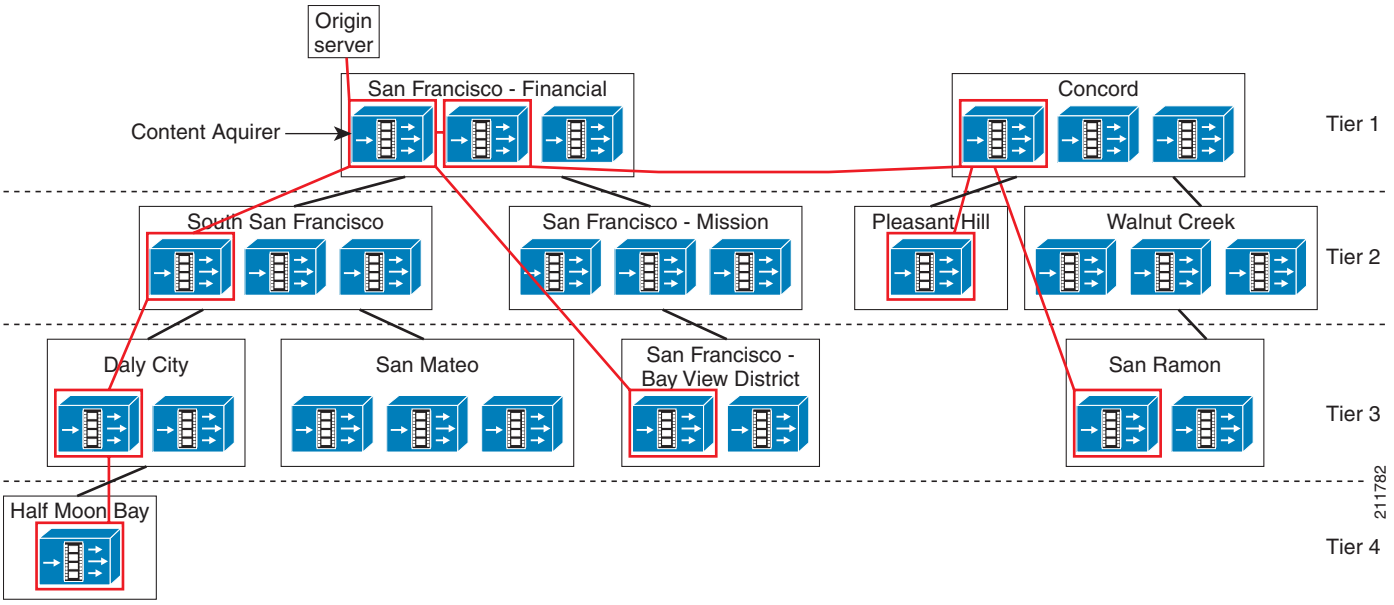

The Cisco CDS supports two types of delivery services:

1. Prefetch/caching delivery services

   For prefetch delivery services, called content delivery services in the CDSM, content is forwarded from Service Engine to Service Engine through the delivery service distribution tree until all Service Engines in the delivery service have received it. The delivery service distribution architecture provides unicast content replication using a hop-by-hop, store-and-forward methodology with the forwarder Service Engines systematically selected on the basis of the manually configured location hierarchy. For caching delivery services, the content need not be fully stored before forwarding.

2. Live delivery service

   The live delivery services are only used for managed live stream splitting. The prefetch/caching delivery services are used for prefetch ingest, dynamic ingest, and hybrid ingest.

There are two methods that can be used to configure a delivery service:

1. Specifying the content by using an externally hosted Manifest file.
2. Specifying the content by using the Internet Streaming CDSM.

   The Internet Streaming CDSM provides a user-friendly interface for adding content and configuring crawl tasks. All entries are validated and a Manifest file is generated. The Internet Streaming CDSM offers the most frequently used parameters, a subset of the Manifest parameters. For a complete set of parameters, use a Manifest file.

The following sections describe the main building blocks of a delivery service:

- Origin Servers
- Manifest File
- Content Acquirer
- Internet Streamer

Origin Servers

Content is stored on origin servers. Each delivery service is configured with one content origin. The same origin server can be used by multiple live delivery services. However, only one prefetch/caching delivery service is allowed per content origin. Each Content Origin is defined in the Internet Streaming CDSM by the following:

- Origin server
- Service routing domain name

The origin server is defined by the domain name that points to the actual origin server. The origin server domain name is used to fetch content that resides outside the delivery service, and to request redirection in case of a failure. The origin server must support at least one of the following protocols in order for the CDS to be able to ingest content:

- HTTP
- HTTPS
- FTP
- CIFS
- SMB

Content can also originate from a local file on the CDS.

The service routing domain name is an FQDN and is used for content redirection. Each content that is ingested by means of the Manifest file is published using the service routing domain name. The service routing domain name configured for the Content Origin must also be configured in the DNS servers, so client requests can be redirected to a Service Router for request mediation and redirection.

Proxy Server

When the Content Acquirer cannot directly access the origin server because the origin server is set up to allow access only by a specified proxy server, a proxy server can be configured. The proxy server is configured through the Internet Streaming CDSM for fetching the Manifest file, and through the Manifest file for fetching the content. Proxy configurations made in the Manifest file take precedence over proxy configurations in the CLI.

Manifest File

The Manifest file contains XML tags, subtags, and attributes used to define how content is ingested and delivered. Each delivery service has one Manifest file. The Manifest file can specify attributes for content playback and control. Attributes for specifying metadata only, without fetching the content, are supported. If special attributes are set, only the metadata and control information are propagated to the Service Engines. The control data is used to control the playback of the content when it gets cached by dynamic ingest. The Manifest file format and details are described in Appendix B, “Creating Manifest Files.”

Crawling

For HTTP, HTTPS, FTP, SMB, or CIFS, a single item can be fetched by specifying a single URL in the CDSM or Manifest file, or content can be fetched by using the crawler feature. The crawler feature methodically and automatically searches acceptable websites and makes a copy of the visited pages for

later processing. The crawler starts with a list of URLs to visit, identifies every web link in the page, and adds every link to the list of URLs to visit. The process ends after one or more of the following conditions are met:

- Links have been followed to a specified depth.
- Maximum number of objects has been acquired.
- Maximum content size has been acquired.

The crawler works as follows:

1. The Content Acquirer requests the starting URL that was configured for the delivery service.
2. The crawler parses the HTML at that URL for links to other files.
3. If links to other files are found, the files are requested.
4. If those files are HTML files, they are also parsed for links to additional files.

In this manner, the Content Acquirer "crawls" through the origin server.

Note The crawler cannot parse JavaScript or VBScript to get the links, nor does it work with HTTP cookies.

A website that has indexing enabled and the default document feature disabled generates HTML that contains a directory listing whenever a directory URL is given. That HTML contains links to the files in that directory. This indexing feature makes it very easy for the crawler to get a full listing of all the content in that directory. The crawler searches the folders rather than parsing the HTML file; therefore, directory indexing must be enabled and the directory cannot contain index.html, default.html, or home.html files.

In FTP acquisition, the crawler crawls the folder hierarchy rather than parsing the HTML file. Content ingest from an SMB server for crawl jobs is similar to FTP ingest; that is, the crawler crawls the folder hierarchy rather than parsing the HTML file.

Content Acquirer

The Content Acquirer parses the Manifest file configured for the delivery service and generates the metadata. If the hybrid ingest attributes are not specified, the Content Acquirer ingests the content after generating the metadata. The Content Acquirer can be shared among many delivery services; in other words, the same Service Engine can perform the Content Acquirer role for another delivery service.

SMB Servers

The CDS supports file acquisition from Windows file servers with shared folders and UNIX servers running the SMB protocol. The Content Acquirer first mounts the share folder. This mount point then acts as the origin server from which the content is fetched. The Content Acquirer fetches the content and stores it locally.

Note With SMB, files greater than two gigabytes cannot be ingested.

HTTP Servers

The no-cache directive in an HTTP server response header tells the client that the content requested is not cacheable. When an HTTP server responds with a no-cache directive, the Content Acquirer behaves as follows:

- If the content to be ingested is specified in an <item> tag in the Manifest file, the Content Acquirer ignores the no-cache directive and fetches the content anyway.
- If the content to be acquired is specified in a <crawler> tag in the Manifest file, the Content Acquirer honors the directive and does not fetch the content.

Internet Streamer

The Internet Streamer application on the Service Engine participates in the delivery service by distributing content within the CDS and delivering content to the clients. The Service Engines can be shared among other delivery services.

HTTP Download—Disabling

Note The ability to disable HTTP downloads on a per-delivery service basis is a feature of Release 2.1.

In some instances, for example when there are contractual obligations to prevent clients from downloading content, it may be necessary to disable HTTP downloads on a delivery service. When HTTP download is disabled, the Web Engine returns a 403 forbidden message.

Service Workflow

What follows is a description of the workflow of a delivery service. Table 2-1 shows sample values for for the delivery service workflow described in Figure 2-3. The delivery service workflow is described in detail following Figure 2-3.

Table 2-1 Delivery Service Parameters Example

Parameter	Value
Type	Caching/Prefetch
Origin Server	www.ivs-internal.com
Service Routing Domain Name	cr-ivs.videonet.com
Delivery Service Contents	http://www.ivs-internal.com/video/wmv-152 http://www.ivs-internal.com/video/wmv-92 http://www.ivs-internal.com/video/wmv-212 http://ww.ivs-internal.com/video/wmv-59 type="cache" http://www.ivs.internal.com/video/wmv-6 type="cache"

Figure 2-3 Delivery Service Workflow Diagram

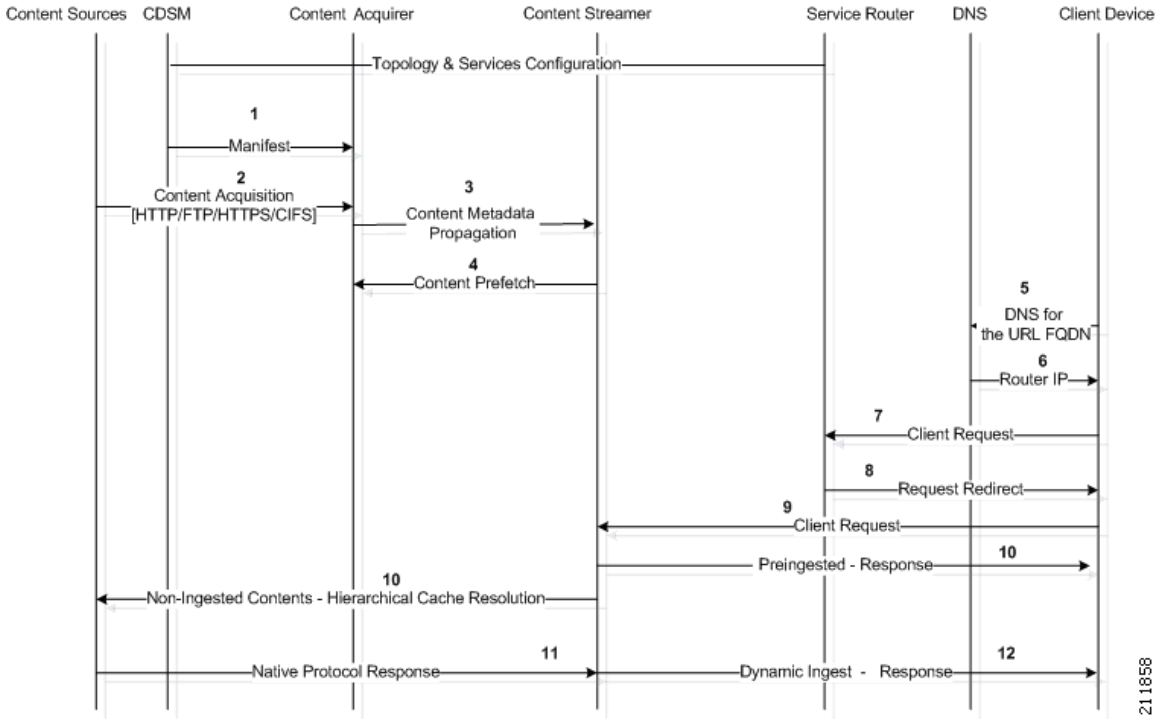

1. The topology is propagated to all the devices registered and activated in the Internet Streaming CDSM. The delivery service configuration is propagated to all the Service Engines subscribed to the delivery service. The Manifest file information is sent to the Content Acquirer for the delivery service.
2. The Content Acquirer parses the Manifest file and generates the metadata. All content listed in the Manifest file, except for non-cache content types, is fetched.
3. The Content Acquirer propagates the metadata to all other Service Engines.
4. The Service Engines receive the metadata and associated prefetched content. The Service Engines do not prefetch content that is “wmt-live” or “cache” types. The “wmt-live” type corresponds to the Windows Media live streaming and the “cache” type corresponds to the hybrid ingest content.
5. The client request for a URL first performs a DNS resolution. The Service Router is configured as the authoritative DNS server for the hosted, or service routing, domain. The URLs that are published to the users have the service routing domain names as the prefix.
6. The Service Router resolves the service routing domain name to its own IP address.
7. The client sends the request to the Service Router and the Service Router uses its routing method to determine the best Service Engine to stream the requested content.
8. The Service Router redirects the client to the best Service Engine.
9. The client sends the request to the Service Engine.

The following are the possible scenarios after the request reaches the Service Engine:

- **Prefetched/Pinned Content**

  Flow 10, “Pre-ingested response.”

  The content is prefetched using the URL: http://www.ivs-internal.com/video/wmv-152

  The actual user request is: http://cr-video.videonet.com/video/wmv-152

The Service Engine processes the user request, and based on the metadata, determines the content was prefetched and pinned in its local storage. The Service Engine looks up the policies for the content and streams the content to the user.

- **Dynamic Ingest/Cached Content**

  Flows 10, 11, 12, "Non-ingested contents—Hierarchical cache resolution," "Native Protocol Response," and "Dynamic ingest response."

  If the request for content is not specified in the Manifest file, dynamic ingest is used.

  The user request is: http://cr-video.videonet.com/video/wmv-cached.wmv

  The Service Engines in the delivery service form a hierarchy, pull the content into the CDS, and cache it. The Service Engine streams the content to the user.

- **Hybrid Ingest/Metadata Only Content**

  (no content flow)

  The request for content is specified in the Manifest file as "cache."

  The user request is: http://cr-video.videonet.com/video/wmv-59

  The Service Engine fetches the content, similar to the dynamic ingest method, but the metadata attributes (for example, serveStartTime, serveStopTime) are honored by the Service Engines and the content is served only if the request falls within the defined time interval.

Programs

A program in the CDS is defined as a scheduled live or rebroadcast event that streams content to client devices. The CDS streams live or rebroadcast content by using the Movie Streamer or the Windows Media Engine.

Movie Streamer live and rebroadcast programs can have multiple tracks (1–3 tracks).

Live Programs

Live events are streamed from third-party encoders (such as Windows Media Encoder Version 9 or the QuickTime encoder) or from streaming servers (such as Windows Media Server). The live stream is ingested by the Content Acquirer and transmitted to all Service Engines using either unicast or multicast. The live stream is transmitted to end users by using either multicast or multicast/unicast live splitting. The live stream is only available to end users during its scheduled times.

With live stream splitting, administrators do not have to create scheduled multicast events, because the Service Engines automatically split the stream.

Unicast to multicast streaming is a solution similar to live stream splitting, except that in the final delivery segment the stream is converted to multicast to minimize the bandwidth demand on the CDS network and to minimize the load on the Service Engines.

Each live program can have up to ten different playtimes scheduled. The program is broadcast from all Service Engines simultaneously.

Rebroadcasts

In a scheduled rebroadcast, prefetched content is scheduled to be streamed from the Service Engines using multicast. Content can only be selected from one delivery service. The Service Engines and device groups assigned to the delivery service are automatically selected when the content files are chosen for the program.

API Program File

Programs can be defined through the Internet Streaming CDSM or through an API. Programs created through APIs are based on a program file. A program file is an XML file that resides on an external server and contains the elements that define the schedule, content, and presentation parameters. The Internet Streaming CDSM gets the program file, parses it, and saves the program file to the database. The program is automatically updated at intervals by refetching the program file and reparsing it. RTSP is the only protocol supported in the program file.

Programs created using an API can be viewed in the Internet Streaming CDSM as read-only, and modifications to the API programs can be accomplished through the API. The API program can also be edited using the Internet Streaming CDSM; however, the information about the API program file is deleted and the program can no longer be modified through the API. A third option is to copy the API program using the Copy Program feature.

CHAPTER 3

Getting Started

This chapter discusses initial device configuration, logging into and navigating the Internet Streaming CDSM, and a typical CDS configuration workflow. This chapter presents the following major topics:

- Initially Configuring the Devices, page 3-1
- Logging In to the Internet Streaming CDSM, page 3-1
- Activating and Synchronizing the Devices, page 3-3
- Navigating the Internet Streaming CDSM, page 3-6
- Configuring Primary and Standby CDSMs, page 3-8
- Typical Configuration Workflow, page 3-10

Initially Configuring the Devices

You must initially configure the Content Delivery Engines (CDEs) before they can participate in the CDS network. The CDE that runs the Internet Streaming CDSM must be initialized first so that the CDEs running the Service Engine (SE) and Service Router (SR) can register with it. For more information about initially configuring the CDEs, see the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide*. See the "Related Publications" section on page xvi.

After you have initially configured your CDEs, you must activate the SEs and SRs and configure the internal clocks by using the Internet Streaming CDSM. See the "Activating and Synchronizing the Devices" section on page 3-3 for more information.

Logging In to the Internet Streaming CDSM

To log in to the Internet Streaming CDSM, do the following:

Step 1 Using your web browser, enter the IP address of your CDSM and port 8443.

For example, if the IP address of your CDSM is 192.168.0.236, enter:

https://192.168.0.236:8443

The Security Alert message is displayed.

Note If you are using Mozilla Firefox version 3.01 or higher as your web browser, you need to add the CDSM IP address to the exception list. After entering the CDSM IP address with port 8443, Firefox displays a Secure Connection Failed message with a link stating "Or you can add an exception." Click this link, then click **Add Exception**. The Add Security Exception dialog box is displayed. Click **Get Certificate**. and then click **Confirm Security Exception**. THE CDSM IP address has been added to the exception list and you will no longer get the Secure Connection Failed message.

Note Sometimes the CDSM is not initially accessible from a web browser. If this occurs you must disable and re-enable the Centralized Management System (CMS). log in to the CLI for the CDSM, and enter the global configuration command **no cms enable** followed by **cms enable**.

Step 2 Click **Yes** to accept the security certificate. The Login page is displayed (Figure 3-1).

Figure 3-1 Internet Streaming CDSM Login Page

Step 3 Enter the username and password and click **Login**. The Internet Streaming CDSM home page is displayed.

The built-in username is *admin* and the initial password is *default*.

Note We strongly recommend that you change the built-in admin password as soon as possible. To do so, log in to the CLI of the CDSM device, and use the **username** *admin* **password** <*password*> global configuration command.

Note If the default username and password have been changed by another CDSM administrator, you need to get the new username and password.

Activating and Synchronizing the Devices

The CDS administrator approves a device by making it active. This security feature prevents unauthorized devices from joining the CDS.

Caution All devices must be synchronized with each other in order for the CDS to function properly.

Synchronization ensures accurate timestamps in all the logs and accuracy in caching decisions determined by If Modified Since (IMS) lookups. Using Network Time Protocol (NTP) to synchronize the devices in the CDS is the best practice.

Note If the network is not configured with NTP, then every device in the CDS must be configured with exactly the same time and time zone. We recommend that you use an NTP server for network synchronization.

Activating and Setting NTP for Each Device

Tip To navigate within the Internet Streaming CDSM, click one of the tabs (for example, Devices) and then one of the tab options (for example Locations). Navigational directions in procedures are written in the following way:
Devices > Devices > Assignments > Device Groups

Note From the Devices Table, you can activate all inactive devices by clicking the **Activate All Inactive SEs** icon. See the "Activating All Inactive Service Engines" section on page 3-5.

To activate and synchronize a Service Engine (SE) or Service Router (SR), do the following:

Step 1 From the Internet Streaming CDSM home page, choose **Devices > Devices**. The Devices Table is displayed (Figure 3-2) listing all the registered SEs and SRs.

Figure 3-2 Devices Table Page–Edit Device

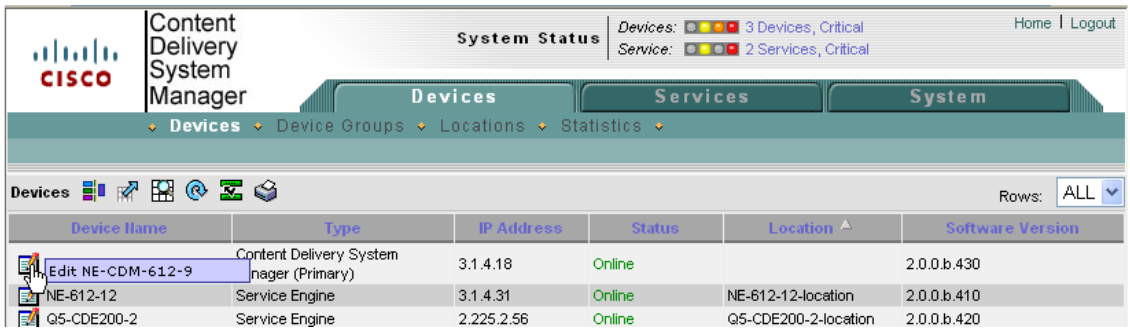

Step 2 Click the **Edit** icon next to the device name. The Device home page is displayed.

Note If the device you want to activate is not listed in the Devices Table, restart the CMS for that device by telneting to it and entering **no cms enable** followed by **cms enable** in global configuration mode.

Step 3 Click **Activate** in the Device home page. The Location dialog box is displayed (Figure 3-3).

Figure 3-3 Device Home Page—Location Dialog Box

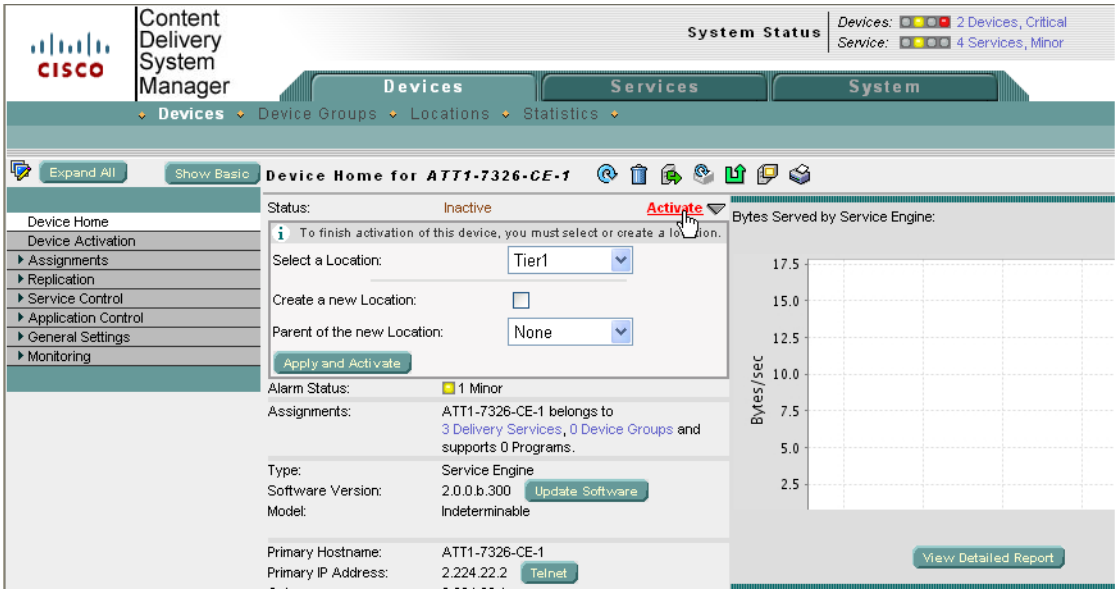

Step 4 Create or choose a location. To activate an SE, you need to assign it to a location.

Because the standby CDSM is global to the CDS network, it does not need to be assigned to a location.

You have the following options in creating or choosing a location:

a. If you have already created locations, you can choose a location from the **Location** drop-down list.

b. To create a default location, which can be edited later, check the **Create a New location** check box. A default location is created with the following name: <*SE-name*>-*location*. From the **Parent of the New Location** drop-down list, choose a parent for this location.

For information about creating locations, see the "Configuring Locations" section on page 4-1.

Step 5 Click **Apply and Activate**.

The Status of the device shows "pending" until the device is fully activated. This may take a few minutes.

Step 6 To display the top-level Table of Contents, click the **Show All** button above the Contents pane.

Step 7 From the left-panel menu, select **General Settings > Network > NTP**. The NTP Settings page is displayed.

Step 8 Check the **Enable** check box and enter the IP address or hostname of each NTP server. Use a space to separate each server.

Step 9 Click **Submit** to save your settings.

The activation and NTP server settings must be completed for each SE, SR, and standby CDSM.

Tip For a quick way to get to other SEs, click the **Display All Devices** icon located to the left of the Expand All button. This icon toggles between the Display All Devices and Menu icons.

For more detailed information about configuring locations, activating devices, and configuring NTP servers, see the following sections:

- Configuring Locations, page 4-1
- Activating a Service Engine, page 4-6
- Configuring NTP, page 4-58

Activating All Inactive Service Engines

To activate all inactive SEs, do the following:

Step 1 From the CDSM home page, choose **Device > Devices** and click the **Activate All Inactive SEs** icon. See Figure 3-4.

Figure 3-4 Devices Table Page—Activate All Inactive Service Engines

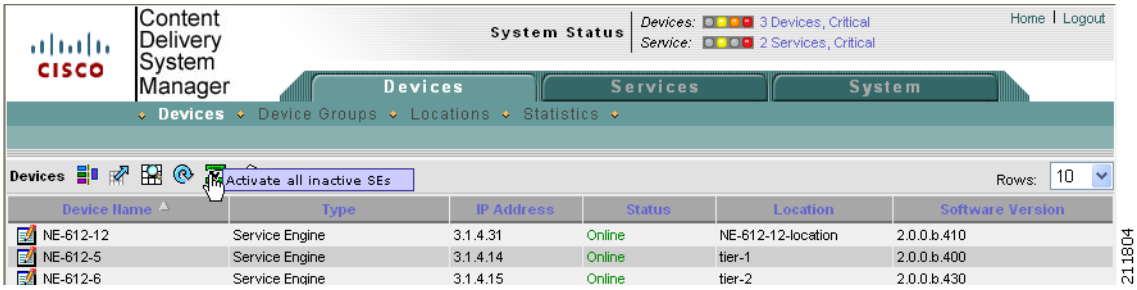

The Location Choice page is displayed (Figure 3-5).

Figure 3-5 Location Choice Page

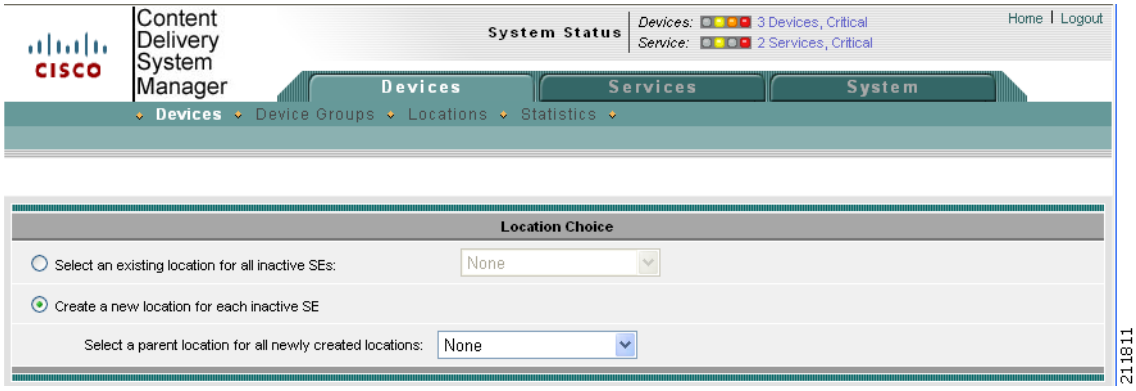

Step 2 In the Location Choice page, click either **Select an Existing Location for All Inactive SEs** or **Create a New Location for Each Inactive SE**.

If you are creating a new location, you can select a parent location, or leave the default of "none."

Step 3 Click **Submit** to save the settings.

The Status in the Devices Table for all the inactive SEs shows "pending" until the devices have been fully activated.

Note All devices activated in this way need to have the NTP settings configured. See Step 6 through Step 9 in the "Activating and Setting NTP for Each Device" section on page 3-3.

Navigating the Internet Streaming CDSM

Figure 3-6 shows the different elements of the Internet Streaming CDSM.

Figure 3-6 Internet Streaming CDSM User Interface

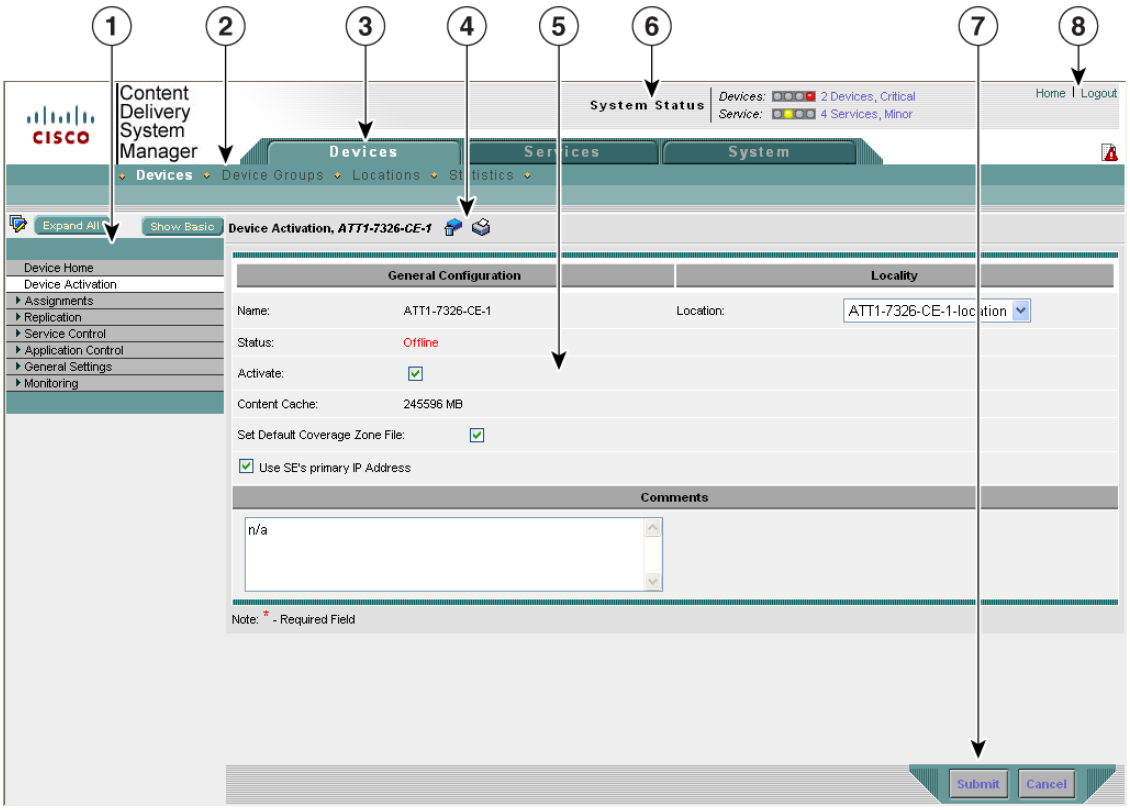

1	Left panel menu	**5**	Page
2	Tab options	**6**	System Status bar
3	Tabs	**7**	Submit and Cancel buttons
4	Task bar	**8**	Tools (Home and Logout)

The System Status bar, tabs, tab options, and tools are accessible from any page in the CDSM. The left panel menu changes depending on which tab and tab option you choose.

Devices, Services, and Other Tables

The Devices Table page shows all the devices registered in the CDSM. Figure 3-7 shows an example of the Devices Table page. A table is displayed for each of the following tab options:

- Devices (from Devices tab)
- Device Groups (from Devices tab)
- Locations (from Devices tab)
- Delivery Services (from Services tab)
- Live Video (from Services tab)

Figure 3-7 Devices Table Page

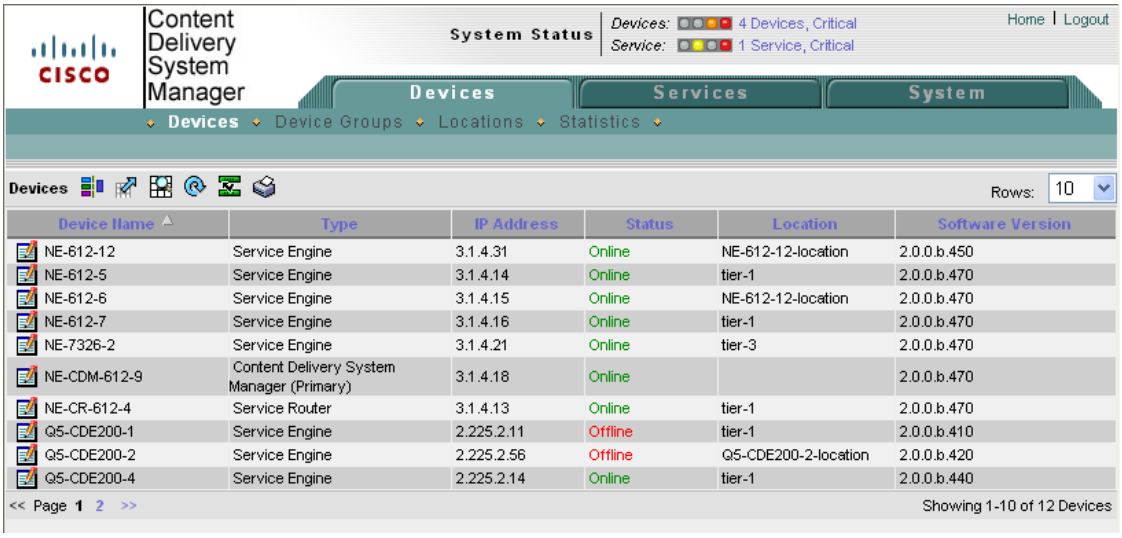

You can sort the information in the table by clicking on any column title. The table can be sorted in ascending or descending order for each column. The task bar options provide other table manipulations, including filtering, refreshing the table, viewing all items, and printing.

The bottom of the table lists the page number and the total number of pages, as well as how many items are showing out of the total number of items.

The table defaults to listing ten rows. You can change the number of rows shown by clicking the Rows drop-down list.

To get more information on an item or to configure an item, click the **Edit** icon to the left of the item name. To create a new item, click the **Create New** icon in the task bar.

Device Home Page

The Device home page provides information about the device, as well as the ability to perform the following tasks:

- Activate the device

- Telnet to the device (In Release 2.3, Telnet is disabled by default.)
- Update the device software
- Assign the device to baseline groups

From the Device home page you can access the delivery services and device groups the device is assigned to, by clicking the appropriate link. All delivery services, or device groups (depending on which link you clicked), configured in your CDS are displayed. Through this page, you can assign the device to additional delivery services or device groups by clicking the icon next to the applicable delivery services or device groups and submitting your selection.

The Device home page offers detailed bandwidth and bytes-served graphs with detailed reports for each.

The left panel menu has two toggle buttons: Show Basic/Show All and Expand All/Collapse All.

- **Show All** Shows all the menu items in the menu.
- **Show Basic** Shows only the Device home menu item.
- **Expand All** Shows every menu and submenu.
- **Collapse All** Shows only the top-level menu items.

Task Bar

The task bar displays information about the page you are on and provides associated tasks. All task bar icons, as well as other icons, have labels that are displayed when you roll over the icon with your mouse pointer.

Any icon used in a procedure is referenced by the rollover label; for example, **Create New** is the rollover label for the following icon:

Configuring Primary and Standby CDSMs

The Internet Streaming CDSM can operate in two different roles: primary and standby. The primary role is the default. You can have only one primary CDSM active in your network; however, you can have any number of CDSMs operating in a standby role to provide redundancy and failover capacity. You must configure the primary CDSM first. See the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide* for information on configuring the primary CDSM.

Note The primary and standby CDSMs must be running the same version of software. You must upgrade your standby CDSM first, and then upgrade your primary CDSM.

To configure a standby CDSM, do the following using the CLI:

Step 1 Follow the instructions for configuring a CDSM using the setup utility, except do not enter the IP address of the CDSM. The instructions can be found in the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide*.

Step 2 Configure the standby CDSM.

```
CDE(config)# cdsm role standby
```

Step 3 Identify the IP address of the primary CDSM.

```
CDE(config)# cdsm ip 10.1.1.90
```

Step 4 Initiate the Centralized Management System (CMS).

```
CDE(config)# cms enable
```

Step 5 Save the configuration.

```
CDE# copy running-config startup-config
```

Step 6 Activate the standby CDSM by using the web interface of the primary CDSM.

The primary CDSM notifies all registered devices that a standby CDSM exists and sends each device the information it needs to contact the standby should the primary fail or become inactive.

Note You cannot log in to the web interface of the standby CDSM. Its function is to maintain an up-to-date copy of the primary's database.

Changing a Standby to a Primary CDSM

Note If your primary CDSM is still operating, you must change its role to standby by executing the **cdsm role standby** command before following these steps. You can only have one primary CDSM operating at any given time.

To change the standby CDSM to become the primary, do the following:

Step 1 If your primary CDSM has failed, enter the following command:

```
CDE(config)# cdsm role primary
```

Step 2 Save the configuration.

```
CDE# copy running-config startup-config
```

Note If you have recently made configuration changes to the primary CDSM, wait at least the polling interval before changing roles to ensure that the standby has a record of the most recent configuration changes.

Typical Configuration Workflow

Once you have completed activating and configuring the NTP servers for all the devices in the CDSM, you are ready to configure the CDS for content delivery. For information about activating and configuring the NTP servers for a device, see the "Activating and Setting NTP for Each Device" section on page 3-3.

Table 3-1 lists the basic tasks for configuring the CDS for content delivery, with references to the associated sections in each chapter.

Table 3-1 Configuration Workflow

Task	Description	Where to Find More Information
Change admin password	Change the administrator password on each device, including the CDSM, and change the administrator password for the system	Log in to the CLI for the device and use the **username** *admin* **password** <*password*> global configuration command. The password strength must be a combination of alphabetic character, at least one number, at least one special character, and at least one uppercase character. To change the password for the CDSM GUI and CLI, go to "Creating, Editing, and Deleting Users," page 6-2
Configure Dedicated Management with Redundant Port	Separate management traffic from application traffic, and configure a redundant port for management	"Configuring Port Channel," page E-1
Create Device Groups	Group like devices to speed up configuration	"Configuring Device Groups," page 4-3
Configure TCP	Configure Device Group TCP settings	"Configuring TCP Settings," page 4-62
Configure RCP	Configure Remote Copy Protocol (RCP) to listen for requests on TCP port 514	"Enabling RCP," page 4-58
Configure FTP	Enable FTP services to listen for connection requests	"Enabling FTP Services," page 4-57
Configure Web Engine	For all SEs participating in delivering content	Begins with the "Configuring Web Engine HTTP Connections," page 4-43
Configure Windows Media Engine	For all SEs participating in delivering Windows Media content	Begins with the "Configuring Windows Media Streaming—General Settings," page 4-32
Configure Movie Streamer	For all SEs participating in delivering MPEG or MOV content	"Configuring Movie Streamer—General Settings," page 4-40
Configure Flash Media Streaming[1]	For all SEs participating in delivering Flash Media Streaming	"Configuring Flash Media Streaming," page 4-42

Table 3-1 Configuration Workflow (continued)

Task	Description	Where to Find More Information
Create Coverage Zone File	Map SEs to client service areas by IP address or geographic location	Appendix C, "Creating Coverage Zone Files"
Import or Upload Coverage Zone File	Apply Coverage Zone mappings to CDS	"Coverage Zone File Registration," page 6-9
Configure Global Routing Method	Set the Coverage Zone file	"Configuring Global Routing" section on page 6-11
Configure Routing Method	Configure the routing method used by SRs	"Configuring the Service Router," page 4-87
Configure Content Origins	Define all origin servers that are used in delivery services.	"Content Origins," page 5-1
Create Delivery Service Definitions	Create delivery services for both prefetched or cached content and live programs	"Creating Delivery Service," page 5-4
Create Live Programs	Create live programs, or rebroadcasts and schedules.	"Configuring Programs," page 5-26

1. Flash Media Streaming is a Release 2.1 feature, therefore this menu option is not available in Release 2.0.

CHAPTER 4

Configuring Devices

This chapter discusses configuring locations and device groups for devices, and detailed instructions on configuring the different types of devices–CDSMs, SEs, and, SRs. This chapter presents the following major topics:

- Configuring Locations, page 4-1
- Configuring Device Groups, page 4-3
- Configuring the Service Engine, page 4-6
- Configuring the Service Router, page 4-87
- Configuring the CDSM, page 4-91

Configuring Locations

Locations are set up in the Internet Streaming CDSM to organize and group SEs into virtual networks for distribution of content through delivery services. For more information about locations, see the "Cisco CDS Topology" section on page 2-1.

Locations need to be configured before you can activate SEs and SRs and bring them online in the CDS network.

To create a new location or edit an existing one, do the following:

Step 1 Choose **Devices > Locations**. The Locations Table page is displayed (Figure 4-1).

Figure 4-1 Locations Table Page

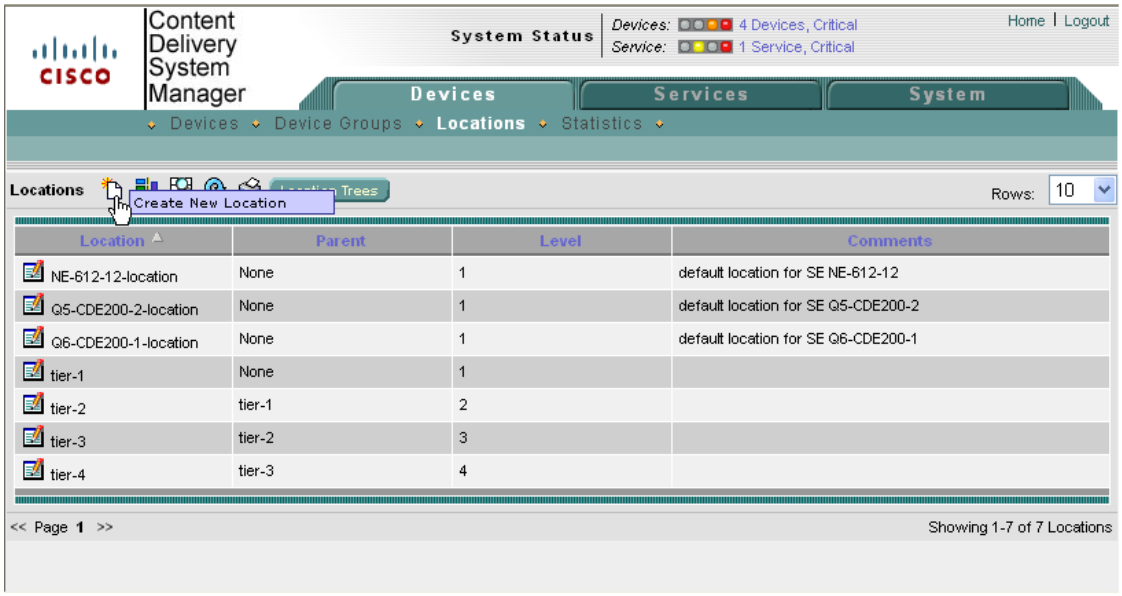

Step 2 In the task bar, click the **Create New Location** icon. The Creating New Location page is displayed (Figure 4-2).

To edit a location, click the **Edit** icon next to the location name.

Figure 4-2 Creating New Location Page

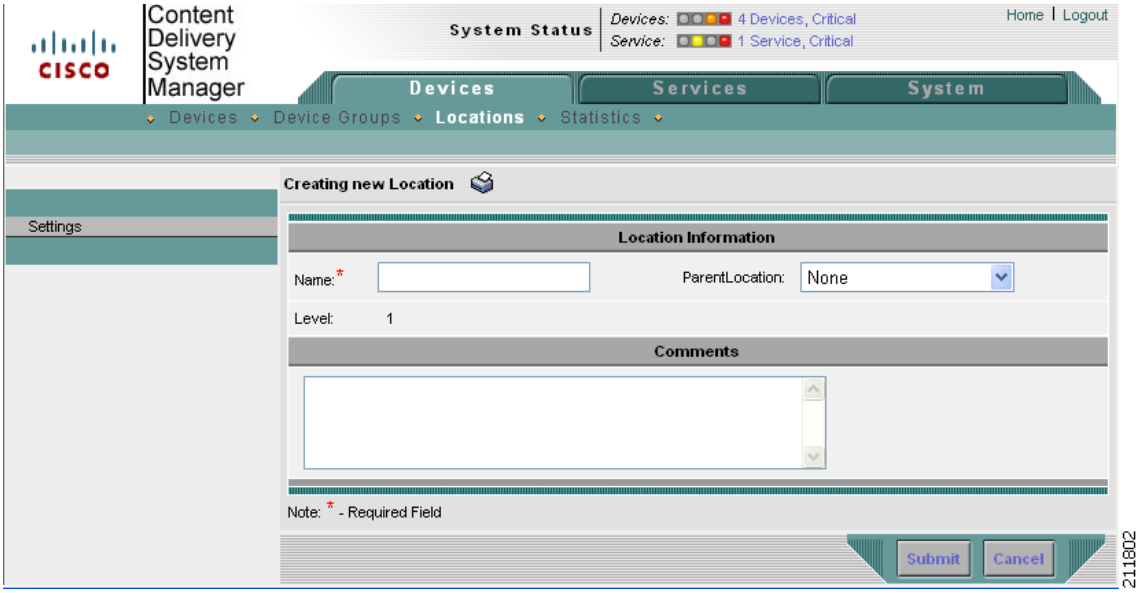

Step 3 Enter the settings as appropriate. See Table 4-1 for a description of the fields.

Table 4-1 Location Fields

Field	Description
Name	Name of the location.
Parent Location	Choose a location from the drop-down list. A location with no parent, None, is level 1. The location level is displayed after you choose a parent location.
Comments	Enter any information about the location.

Step 4 Click **Submit** to save the settings.

To delete a location, from the Locations Table page, click the **Edit** icon next to the location you want to delete, and click the **Trash** icon in the task bar.

To view the location tree, click the **Location Trees** icon in the task bar. The location tree represents the network topology you configured when you assigned a parent to each location.

Configuring Device Groups

The Internet Streaming CDSM allows you to configure SEs into device groups so that the entire group of SEs is configured at one time. Device groups and SEs share the same configuration features and options.

This section covers creating, editing, and deleting device groups. For information on assigning devices to a device group, see the "Assigning Devices to Device Groups" section on page 4-9. All other configuration pages for a device group are covered in the "Configuring the Service Engine" section on page 4-6.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

To create or edit a device group, do the following:

Step 1 Choose **Devices > Device Groups**. The Device Groups Table page is displayed (Figure 4-3).

Figure 4-3 Device Groups Table Page

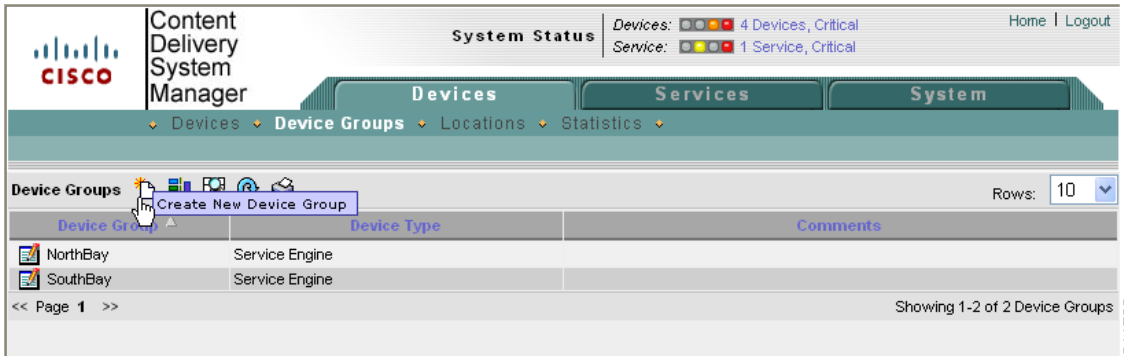

Step 2 In the task bar, click the **Create New Device Group** icon. The Creating New Device Group page is displayed (Figure 4-4).

To edit a device group, click the **Edit** icon next to the device group name.

Figure 4-4 Creating New Device Group Page

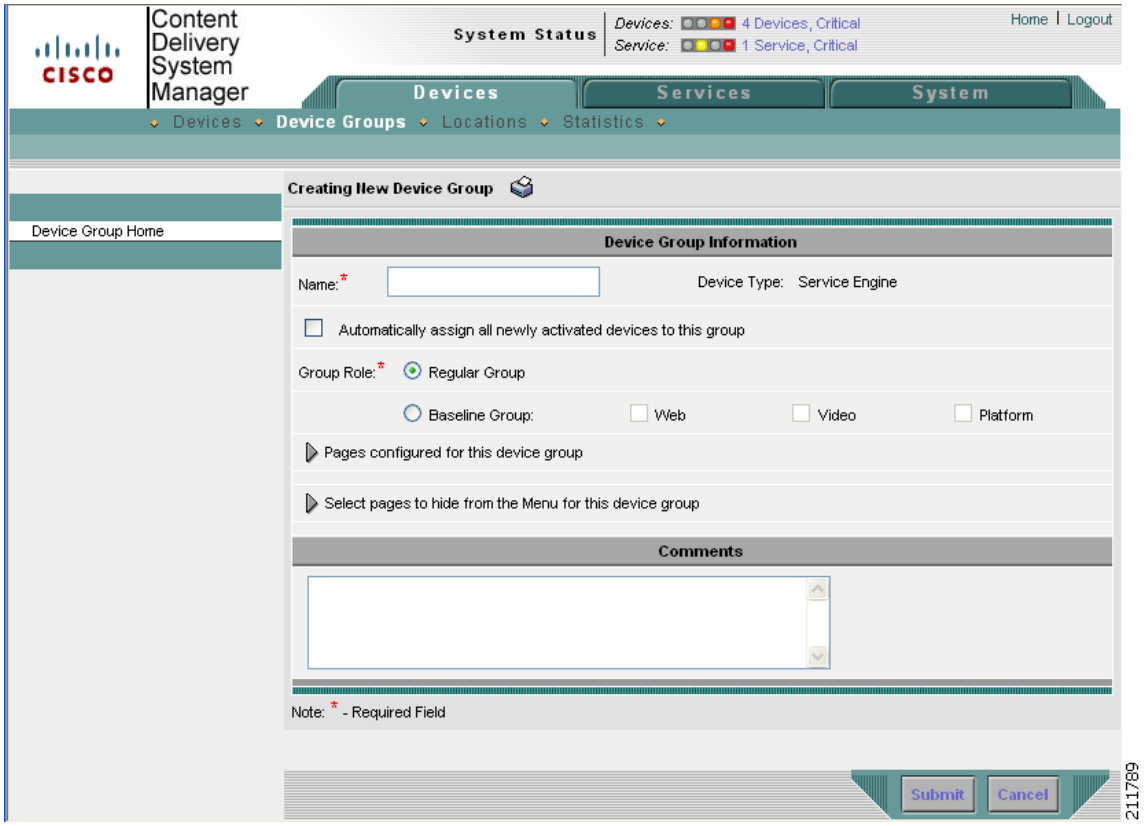

Step 3 In the Name field, enter the name of the device group. The name must be unique and should be a name that is useful in distinguishing the device group from the others in the CDS.

Step 4 Check the **Automatically assign all newly activated devices to this group** check box if applicable.

Step 5 Choose **Regular Group** to indicate this group is not used as a baseline for all SEs or choose **Baseline Group** and select the baseline type to define this group as a baseline for all SEs.

For information about baseline groups, see the "Baseline Groups" section on page 2-3.

Step 6 To customize the left panel menu for this device group, click the **Select pages to hide from the menu for this device group** arrow, and check the pages you want to hide. To collapse these settings, click the arrow again.

Use this feature to remove from view any configuration pages that you do not need for the device group.

Step 7 In the Comments field, enter any information about the device group.

Step 8 Click **Submit** to save the settings.

If you are editing this device group, you can view a list of all settings configured for this device group by clicking the **Pages configured for this device group** arrow. To collapse this information list, click the arrow again.

To delete a device group, click the **Trash** icon in the task bar.

Step 9 To assign SEs to the device group, choose **Assignments > Devices**. The Assignment table is displayed listing all SEs in the CDS.

Step 10 Click the **Assign** icon (blue cross mark) next to each SE name you want to assign to this group.

To assign all SEs, click **Assign all Service Engines** in the task bar.

Step 11 Click **Submit** to add the selected SEs to the device group.

To remove an SE from the device group, click the **Unassign** icon (green check mark) next to the name of the SE, and click **Submit**.

To remove all SEs from the device group, click the **Unassign all Service Engines** icon in the task bar, and click **Submit**.

Note All configuration settings for a device group can be configured on an individual SE as well. All other configuration pages for a device group are covered in the "Configuring the Service Engine" section on page 4-6.

Device Group Overlap

If you want the ability to assign a device to more than one device group, you must enable device group overlap. Device group overlap is enabled by default.

To enable or disable device group overlap, do the following:

Step 1 Choose **System > Configuration**. The Config Properties page is displayed.

Step 2 Click the **Edit** icon next to the DeviceGroup.overlap property. The Modifying Config Property page is displayed.

Step 3 To enable device group overlap, choose **true** from the **Value** drop-down list.

To disable device group overlap, choose **false** from the **Value** drop-down list.

Step 4 Click **Submit** to save the settings.

If you disable device group overlap after you have assigned devices to multiple device groups, existing overlaps are maintained. Any newly added groups will not allow assignment of devices that are already assigned to another group, and new devices cannot be added to any groups with overlapping devices.

Tip To force the complete configuration set of a device group to all devices in that group, click the **Force Group Settings** icon in the task bar.

Configuring the Service Engine

This section walks you through the different configuration pages available for a Service Engine. The main configuration groups are described as follows:

- Service Control—Settings for access control by way of client request filtering, URL signing, and third-party QoS and conditional access policies; additionally, transaction logs are configured to monitor traffic
- Application Control—Settings for bandwidth management of delivery services and protocol engines (Web, Windows Media, Movie Streamer, Flash Media Streaming, and RTSP advanced settings)
- General Settings—Settings for access control of the device, maintenance, network connectivity, and monitoring

The first two pages, Device Activation and Assignment, cover activating an SE in the Internet Streaming CDSM and assigning it to a location, and assigning device groups to the SE.

Note All SE settings in this section, except those listed below, can also be configured for a device group by choosing **Devices > Device Groups**. The following pages are not available for device group configuration:

- **Devices > Application Control > Windows Media Streaming > Bypass List**. See the "Configuring Windows Media Streaming—Bypass List" section on page 4-39 for more information.
- **Devices > General Settings > Network > Network Interfaces**. See the "Viewing Network Interfaces" section on page 4-61 for more information.
- **Devices > General Settings > Network > External IP**. See the "Configuring External IP Addresses" section on page 4-61 for more information.

Activating a Service Engine

Activating a device (Service Engine, Service Router, or CDSM) can be done through the Device home page initially, or through the Device Activation page.

To activate a device from the Device Activation page, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed (Figure 4-5).

Figure 4-5 Devices Table Page

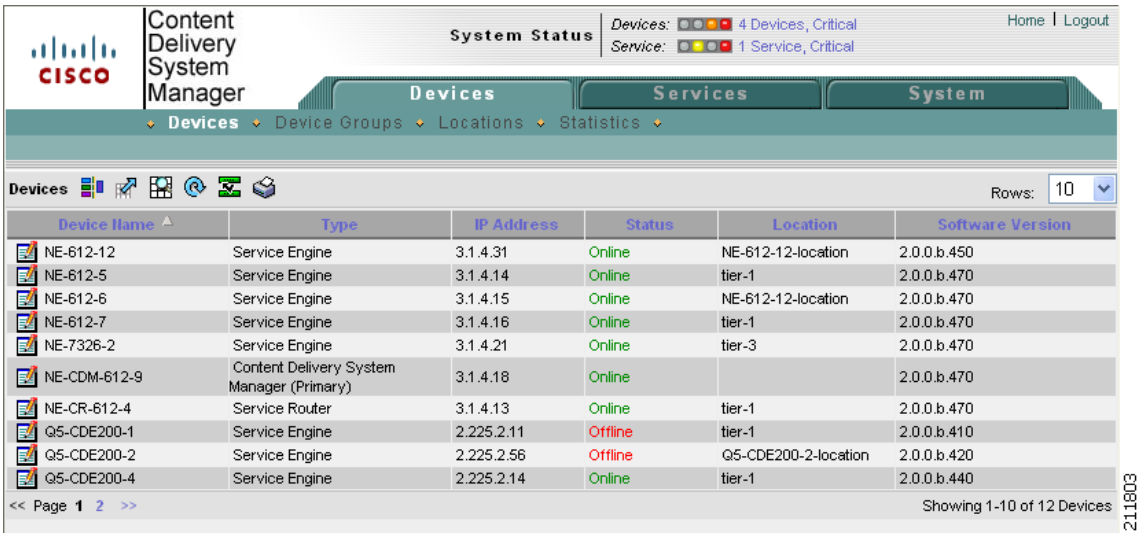

Step 2 Click the **Edit** icon next to the device you want to configure. The Device home page is displayed.

Step 3 Click **Show All** to display the top-level menu options, and click **Device Activation**. The Device Activation page is displayed (Figure 4-6).

Figure 4-6 Device Activation Page

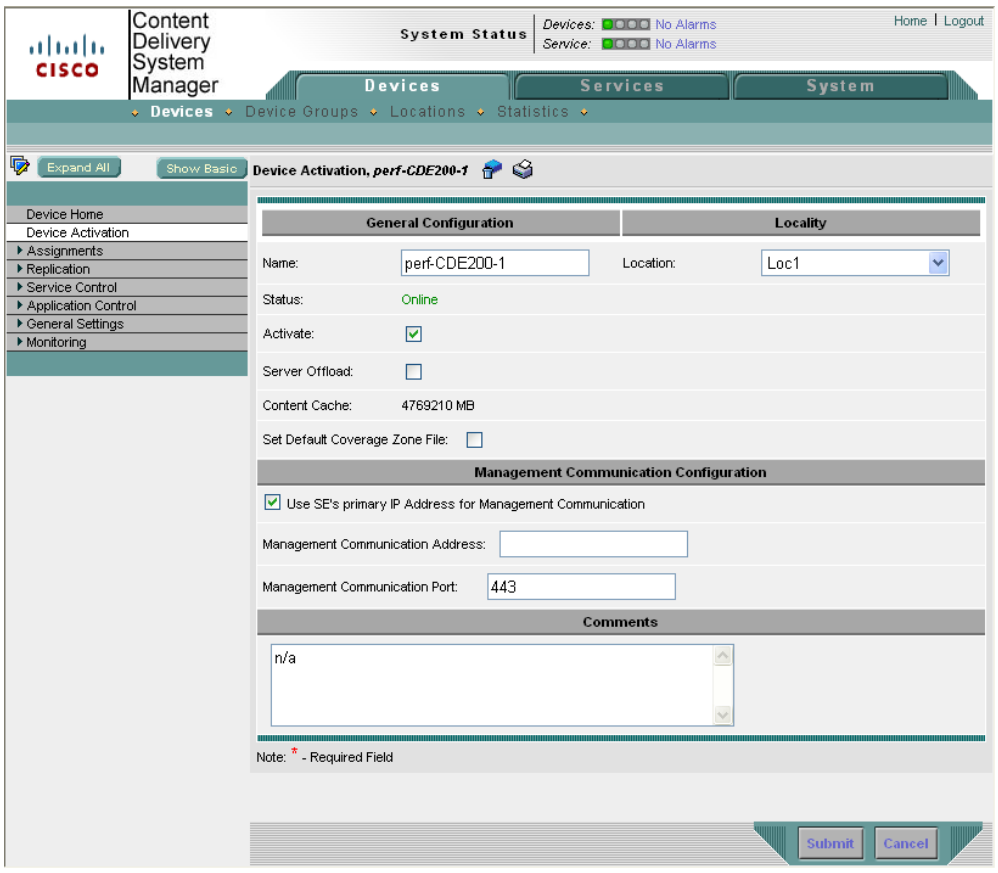

Step 4 Enter the settings as appropriate. See Table 4-2 for a description of the fields.

Table 4-2 Device Activation Fields

Field	Description
Name	Name of the device.
Activate	To activate or deactivate the device, check or uncheck the **Activate** check box.
Server Offload	To offload this device for maintenance or a software upgrade, check the **Server Offload** check box. When checked, the Service Router stops sending requests to this device. **Note** If a client paused a program at that moment **Server Offload** is enabled, most likely resuming the program will fail. To monitor the current streams on an SE during the Server Offload state, use the **show interface** command. If the packets received or packets sent is increasing then the SE is streaming. Packets received will be high if there is an incoming stream. **Note** We recommend separating the management traffic from the streaming traffic by using the port channel configuration, see the "Configuring Port Channel" section on page E-1 for more information. 1. If management and streaming traffic are separated, the **show interface** command for the streaming port channel displays information on active sessions. 2. If management and streaming traffic are not separated, the **show interface** command shows very low traffic; the packets received and packets sent are lower than a client streaming session. Once the SE has finished streaming, you can perform maintenance or upgrade the software on the device. For information about upgrading the software, see the "Upgrading the Software" section on page 8-6. The Status field on the Device Activation page and the Devices Table page displays "offloading" when **Server Offload** is checked. Once the software upgrade or maintenance is complete, you need to uncheck the **Server Offload** check box so that the device can again participate in the system. **Note** If the Server Offload option is set on an SE that is acting as the Content Acquirer for a delivery service for dynamic ingest or live stream splitting, a new SE is chosen as the Location Leader for the delivery service. However, if the Content Acquirer is up and communicating with the CDSM, it continues to perform content ingest and content distribution.
Content Cache	Informational only. The content cache size is the total disk space on the CDS network file system (CDNFS) on the SE that is designated for cache. The Content Cache represents the unused cache space. The used cache space is the disk space allotted for all the delivery services to which the SE is assigned. To view the used cache space, choose **Services > Service Definition > Delivery Services > Assign Service Engines**.

Table 4-2 Device Activation Fields (continued)

Field	Description
Set Default Coverage Zone File	When checked, which is the default setting, a default Coverage Zone file is generated with the SE serving the local subnet it resides on. The coverage zone is a CDS network-wide mapping of client IP addresses to SE IP addresses that should respond to client requests. For more information, see the "Coverage Zone File Registration" section on page 6-9. The default coverage zone can be disabled and you can create and assign custom coverage zones using the Coverage Zone file import or upload. Uncheck the **Set Default Coverage Zone File** check box to use a user-defined Coverage Zone file that was imported or uploaded.
Location	Lists all the locations configured for the CDS.
Use SE's primary IP address	Enables the CDSM to use the IP address on the primary interface of the SE for management communications. **Note** If the **Use SE's primary IP Address for Management Communication** check box is checked and the Management Communication Address and Port are configured, the CDSM uses the SE's primary IP address for communication. If communication attempts to the primary address fail, the CDSM tries the configured IP address and port. **Note** Do not check the **Use SE's primary IP Address for Management Communication** check box if you want to separate management and streaming traffic. Instead, use the Management Communication Address and Port fields to specify where management traffic should be sent.
Management Communication Address	Manually configures a management IP address for the CDSM to communicate with the SE. **Note** This is a Release 2.1 feature. Manual configuration of the management IP address and port are used when using port channel configuration to separate management and streaming traffic. For more information about port channel configuration see the "Configuring Port Channel and Load Balancing Settings" section on page 4-62 and the "Configuring Port Channel" section on page E-1.
Management Communication Port	Port number to enable communication between the CDSM and the SE. **Note** This is a Release 2.1 feature.
Comments	Information about the settings.

Step 5 Click **Submit** to save the settings.

Assigning Devices to Device Groups

You can assign devices to device groups in three ways:

- Through the Device Group Assignment page

- Through the device Assignment page
- Through the Device home page, if the device group is a baseline group

To assign devices to device groups through the Assignment page, do the following:

Step 1 Choose **Devices > Devices**, and click the **Edit** icon next to the device you want to assign.

Step 2 Click **Show All**, and then click **Assignments > Device Groups**. The Assignment page is displayed (Figure 4-7).

Figure 4-7 Assignment Page

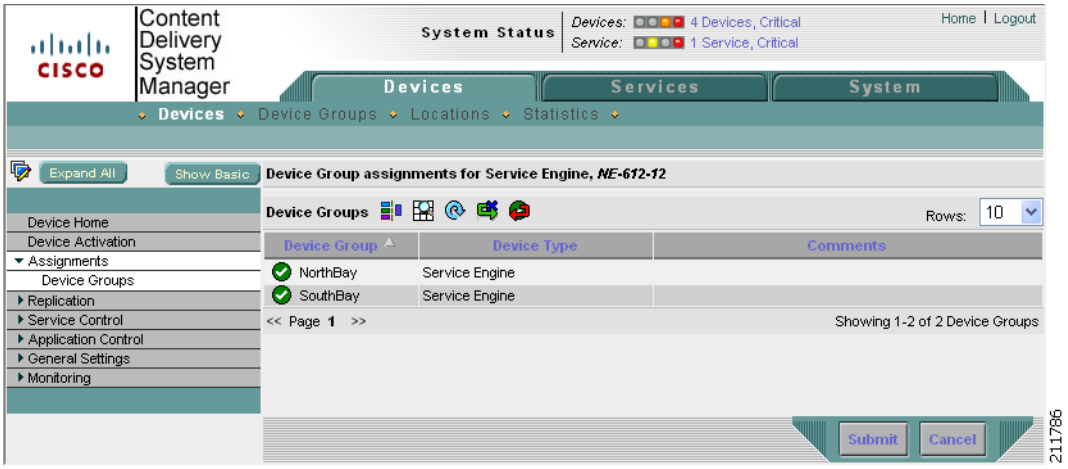

Step 3 Choose **Assignments > Device Groups**. The Device Group Table page is displayed with all of the configured device groups listed.

Step 4 Click the **Assign** icon (blue cross mark) next to the device group you want to assign to this SE. Alternatively, click the **Assign All Device Groups** icon in the task bar.

A green arrow wrapped around the blue X indicates an SE assignment is ready to be submitted. To unassign an SE, click this icon.

Step 5 Click **Submit** to save the settings.

A green circle with a check mark indicates a device group is assigned to this SE. To unassign the device group, click this icon, or click the **Unassign All Device Groups** icon in the task bar. Click **Submit** to save the changes.

Note From this point forward, the beginning steps in the procedures are combined into one step using notation similar to the following: **Devices > Devices Assignments > Device Groups**.

Configuring Bandwidth for Replication and Ingest

The bandwidth used for replication and ingest is determined by the settings in the Default Bandwidth and the Scheduled Bandwidth pages. The replication configuration pages consist of the following:

- Default Bandwidth

- Scheduled Bandwidth

Default Bandwidth

The default bandwidth settings can be configured for acquisition (ingest) and distribution (replication) of content. The default settings are used unless a scheduled bandwidth is configured for a specified time period.

To set the default bandwidth for replication, do the following:

Step 1 Choose **Devices > Devices > Replication > Default Bandwidth**. The Replication Default Bandwidth page is displayed (Figure 4-8).

Figure 4-8 Replication Default Bandwidth Page

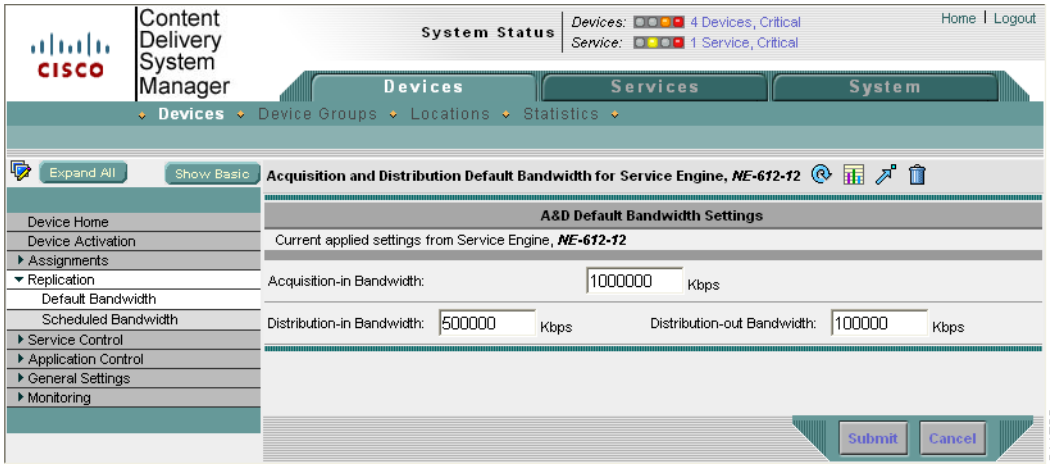

Step 2 Enter the settings as appropriate. See Table 4-3 for a description of the fields.

Table 4-3 Replication Default Bandwidth Fields

Field	Description
Acquisition-in Bandwidth	Bandwidth used for ingesting content when this SE is acting as the Content Acquirer. The default is 500000 kbps (kilobits per second) .
Distribution-in Bandwidth	Bandwidth used for incoming content that is sent by a forwarding SE as part of the distribution process. The default is 250000 kbps.
Distribution-out Bandwidth	Bandwidth used for outgoing content that is sent to a downstream SE as part of the distribution process. The default is 500000 kbps.

Step 3 Click **Submit** to save the settings.

Bandwidth Graph

To view a graphical representation of the bandwidth settings, click the **Display Graph** icon in the task bar. The Acquisition and Distribution Bandwidth graph is displayed in a new window.

The vertical axis of the graph represents the amount of bandwidth in Kbps (kilobits per second) and the horizontal axis represents the days of the week. The scale shown on the vertical axis is determined dynamically based on the bandwidth rate for a particular type of bandwidth and is incremented appropriately. The scale shown on the horizontal axis for each day is incremented for each hour. Each type of bandwidth is represented by a unique color. A legend at the bottom of the graph maps the colors to the corresponding bandwidths.

You can change the graph view by choosing the different options, as described in Table 4-4.

Table 4-4 Acquisition and Distribution Bandwidth Graph—Viewing Options

Option	Description
Distribution In	Bandwidth settings for incoming content distribution traffic.
Distribution Out	Bandwidth settings for outgoing content distribution traffic.
Acquisition In	Bandwidth settings for incoming content acquisition traffic.
All Servers	A consolidated view of all configured bandwidth types. This is the default.
Show Detailed Bandwidth/Show Effective Bandwidth	Toggles between the two options: Show Detailed Bandwidth—Displays detailed bandwidth settings for the SE and its associated device groups. The bandwidth settings of the device and device groups are shown in different colors for easy identification. Show Effective Bandwidth—Displays the composite (aggregate) bandwidth settings for the SE and its associated device groups.
Show Aggregate View/Show Non-Aggregate View	Toggles between the two options: Show Aggregate View—Displays the bandwidth settings configured for the corresponding device groups. Show non-Aggregate View—Displays the bandwidth settings configured for the SE.
Sun, Mon, Tues, Wed, Thurs, Fri, Sat	Displays the bandwidth settings for the corresponding day of the week.
Full Week	Displays the bandwidth settings for the entire week This is the default view and is combined with the All Servers view.

Scheduled Bandwidth

Scheduled Bandwidth settings take precedence over Default Bandwidth settings.

To configure a bandwidth schedule, do the following:

Step 1 Choose **Devices > Devices > Replication > Scheduled Bandwidth**. The Replication Scheduled Bandwidth Table page is displayed (Figure 4-9).

Figure 4-9 Replication Scheduled Bandwidth Table Page

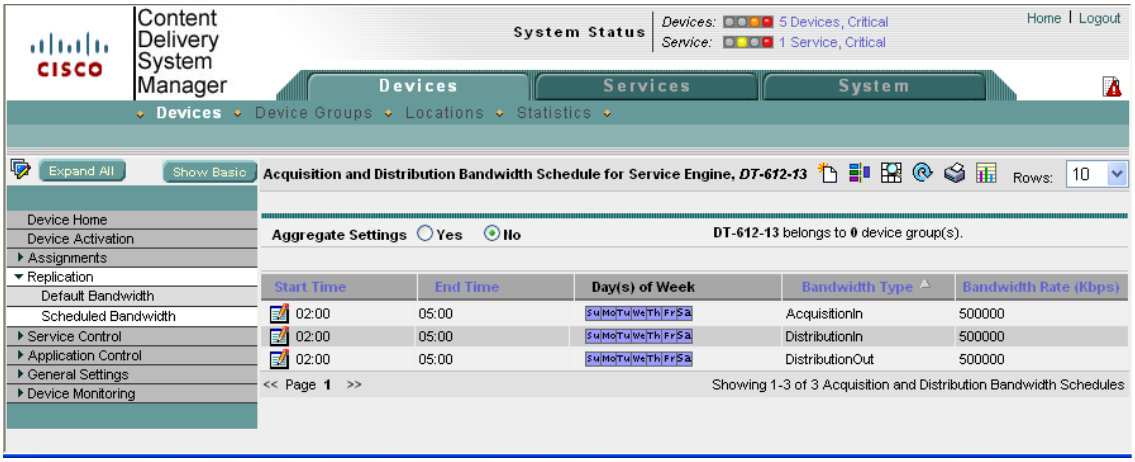

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click the **Create New** icon in the task bar. The Replication Scheduled Bandwidth page is displayed (Figure 4-10).

To edit a scheduled bandwidth, click the **Edit** icon next to the scheduled bandwidth you want to edit.

Figure 4-10 Replication Scheduled Bandwidth Page

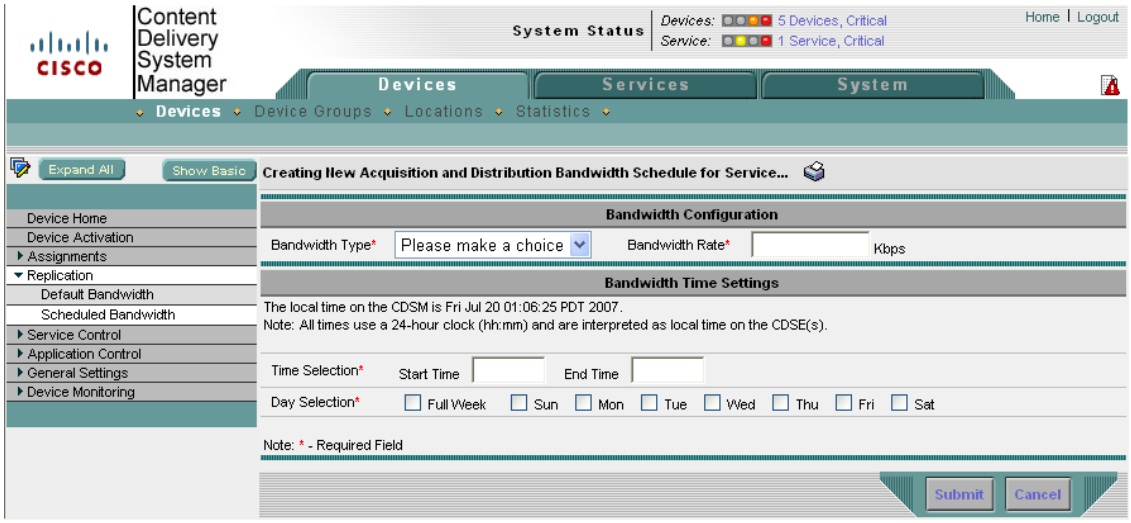

Step 4 Enter the settings as appropriate. See Table 4-5 for a description of the fields.

Table 4-5 Replication Scheduled Bandwidth Fields

Field	Description
Bandwidth Type	Distribution-in—For incoming content distribution traffic from SEs. Distribution-out—For outgoing content distribution traffic to SEs. Acquisition-in—For incoming content acquisition traffic from origin servers.
Bandwidth Rate	Maximum amount of bandwidth that you want to allow (in kbps).
Start Time	Time of day for the bandwidth setting to begin, using a 24-hour clock in local time (hh:mm).
End Time	Time of day for the bandwidth setting to end (hh:mm).
Day Selection	Days on which bandwidth settings apply. • Full Week—Specifies that the allowable bandwidth settings are applied for an entire week. • Sun, Mon, Tue, Wed, Thu, Fri, and Sat—Specifies individual days of the week on which the allowable bandwidth settings take effect.

Step 5 Click **Submit** to save the settings.

Service Control

The Service Control pages provide settings for client request filtering, URL signing, and third-party QoS and conditional access policies. Additionally, transaction logs that monitor traffic are configured under the Service Control. Configuring service control consists of the following procedures:

- Configuring Service Rules
- Configuring ICAP

- Configuring PCMM QoS Policy
- Configuring URL Signing
- Configuring Transaction Logs

Configuring Service Rules

Note This is a licensed feature. Please ensure that you have purchased a Service Rule license for this advanced feature.

The Rules Template licensed feature provides a flexible mechanism to specify configurable caching requests by allowing these requests to be *matched* against an arbitrary number of parameters, with an arbitrary number of *policies* applied against the matches. You can specify a set of rules, each clearly identified by an action and a pattern. Subsequently, for every incoming request, if a pattern for a rule matches the given request, the corresponding action for that rule is taken.

Note The processing time on the SE is directly related to the number of service rules configured. Processing times increase with an increase in the total number of rules configured. If the SE processing time is greater than twice the datafeed poll rate, then the device goes offline until the processing is completed. You can avoid this by configuring a higher datafeed poll rate. The recommended datafeed poll rate for 750 service rules is 300 seconds. To configure the datafeed poll rate, see the "Configuring System Settings" section on page 6-6.

Configuring a service rule consists of the following tasks:

- Enabling the service rules. (Only needs to be performed once.)
- Configuring a pattern list and adding a pattern to it.
- Associating an action with an existing pattern list.

To configure or edit service rule settings, do the following:

Step 1 Choose **Devices > Devices > Service Control > Enable Rules**. The Enable Service Rules page is displayed.

Step 2 Check the **Enable** check box to enable the use of rule settings.

Step 3 Click **Submit** to save the settings.

Step 4 Choose **Devices > Devices > Service Control > Service Rules**. The Service Rules Table page is displayed.

Step 5 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 6 Click the **Create New** icon in the task bar. The Service Rules page is displayed (Figure 4-11).

To edit a service rule, click the **Edit** icon next to the service rule you want to edit.

Figure 4-11 Service Rules Page

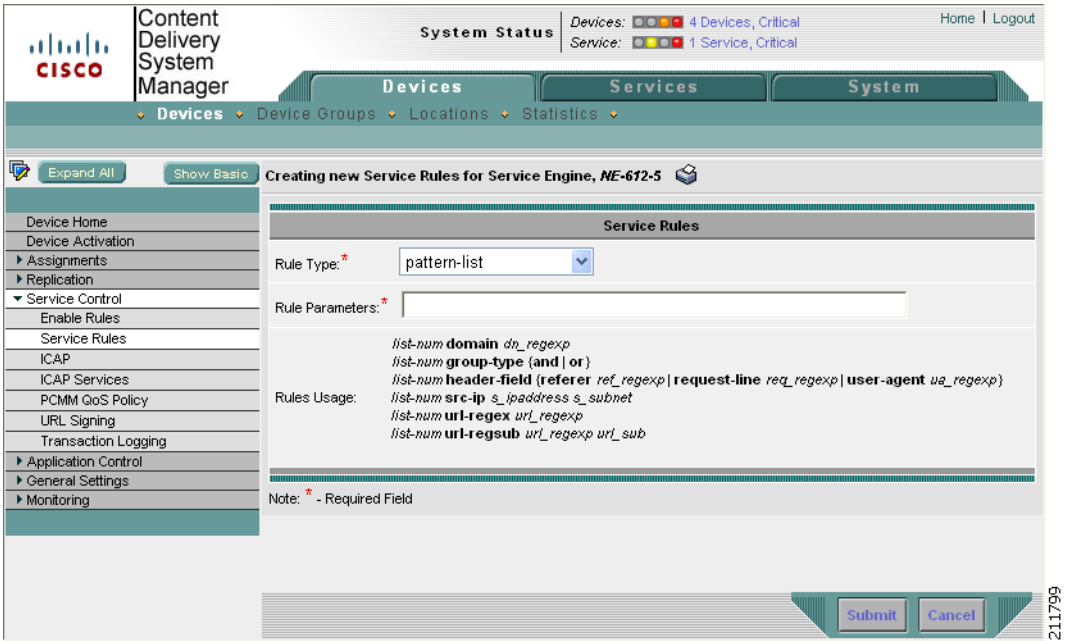

Step 7 Create a pattern list and add a pattern to it.

a. From the **Rule Type** drop-down list, choose *pattern-list*.

b. In the **Rule Parameters** field, configure the pattern list number and the pattern type, following the rules usage guidelines shown on the Service Rules page. See Table 4-6 for a description of pattern types. The rule patterns are not case-sensitive.

For example, to create pattern list number 72 with the pattern type *domain* and the yahoo.com domain as the domain to be acted on, enter **72 domain yahoo.com** in the Rule Parameters field.

Table 4-6 Service Rules Pattern Types

Pattern Type	Description	Syntax
domain	Matches the domain name in the URL or the host header against a regular expression. For example, ".*ibm.*" matches any domain name that contains the "ibm" substring. "\.foo\.com$" matches any domain name that ends with the ".foo.com" substring. In regular expression syntax, the dollar sign ($) metacharacter directs that a match is made only when the pattern is found at the end of a line.	**rule pattern-list** *list_num* **domain** *dn_regexp*
group-type	Patterns can be combined by using the AND or OR function with the group-type pattern (for example, **rule pattern-list 1group-type and**). The default is OR.	**rule pattern-list** *list-num* **group-type** {**and** | **or**}
header-field	Request header field pattern. Request header field patterns **referer**, **request-line,** and **user-agent** are supported for the allow, block, and redirect actions. The **referer** pattern is matched against the Referer header in the request, the **request-line** pattern is matched against the first line of the request, and the **user-agent** pattern is matched against the User-Agent header in the request. **Note** In Release 2.3, Flash Media Streaming supports the **referer** header field pattern for the allow and block actions.	**rule pattern-list** *list_num* **header-field** {**referer** *ref_regexp* | **request-line** *req_regexp* | **user-agent** *ua_regexp*}
scr-ip	Matches the request's source IP address and netmask.	**rule pattern-list** *list_num* **src-ip** *s_ipaddress s_subnet*
url-regex	Matches the URL against a regular expression. The match is not case sensitive.	**rule pattern-list** *list_num* **url-regex** *url_regexp*
url-regsub	For the **rewrite** and **redirect** actions, matches the URL against a regular expression to form a new URL in accordance with the pattern substitution specification. The match is not case sensitive. The valid substitution index range is from 1 to 9.	**rule pattern-list** *list_num* **url-regsub** *url_regexp url_sub*

Note A domain pattern list matching an SE IP address is not supported when IP-based redirection is enabled on the Service Router. Flash Media Streaming bypasses the rules configuration if the request is from another SE. See the "Configuring the Service Router" section on page 4-87 for more information about IP-based redirection.

Step 8 Click **Submit** to save the settings.

The maximum number of pattern lists allowed is 128.

Step 9 Associate an action with an existing pattern list.

a. Choose an action type from the Rule Type drop-down list. See Table 4-7 for a description of rule actions.

b. In the Rule Parameters field, enter the list number of the pattern list that you want to associate with this action.

For example, if you want to block access by any protocol to yahoo.com, then choose **block** from the Rule Type drop-down list, and enter **pattern-list 72 protocol all** in the Rule Parameters field.

Note Movie Streamer and Flash Media Streaming support URL signing in Release 2.2. Flash Media Streaming only supports the following actions: allow, block, and validate-url-signature.

Table 4-7 Service Rule Actions

Action Type	Description	Syntax
allow	Allows incoming requests that match the pattern list. This rule action can be used in combination with block actions to allow selective types of requests. The allow action does not carry any meaning as a standalone action.	**rule action allow pattern-list** *list_num* [**protocol** {**all** | **http** | **rtsp**}]
block	Blocks this request and allows all others.	**rule action block pattern-list** *list_num* [**protocol** {**all** | **http** | **rtsp**}]
generate-url-signature	Generates the URL signatures in the Windows Media metafile response associated with prefetched content, based on the SE configuration for the URL signature and this rule action.	**rule action generate-url-signature** [**include-client-src-ip**] **key-id-owner** *owner_num* **key-id-number** *id_num* **pattern-list** *list_num* [**protocol** {**all** | **http**}]
no-cache	Does not cache this object.	**rule action no-cache pattern-list** *list_num* [**protocol** {**all** | **http**}]
redirect	Redirects the original request to a specified URL. Redirect is relevant to the RADIUS server only if the RADIUS server has been configured for redirect.	**rule action redirect** *url* **pattern-list** *list_num* [**protocol** {**all** | **http** | **rtsp**}]
refresh	For a cache hit, forces an object freshness check with the server.	**rule action refresh pattern-list** *list_num* [**protocol** {**all** | **http**}]
rewrite	Rewrites the original request as a specified URL.	**rule action rewrite pattern-list** *list_num* [**protocol** {**all** | **http** | **rtsp**}]
use-icap-service	Applies ICAP processing and uses a specific ICAP service only for those requests that match this Rules Template action. An ICAP service is a collection of attributes that defines the type of modification to be performed on HTTP requests and responses. If this action is configured, you can allow requests and responses to be processed by ICAP servers for content adaptation.	**rule action use-icap-service** *service-name* **pattern-list** *list_num* [**protocol** {**all** | **http**}]
validate-url-signature	Validates the URL signature for a request using the configuration on your SE for the URL signature and allows the request processing to proceed for this request.	**rule action validate-url-signature error-redirect-url** *url* **pattern-list** *list_num* [**protocol** {**all** | **http** | **rtsp**}]

Step 10 Click **Submit** to save the settings.

Execution Order of Rule Actions

The order in which the rule actions are executed is as follows:

1. block or allow

   Note The allow and block actions carry the same precedence. The order of execution depends on the order of configuration between allow and block actions. Other actions always take precedence over allow.

2. redirect (before cache lookup)
3. rewrite (before cache lookup)
4. refresh (after cache lookup, in the case of cache hit)
5. no-cache
6. use-icap-service
7. generate-url-signature
8. validate-url-signature
9. allow

Note When configuring service rules, you must configure the same service rules on all SEs participating in a delivery service in order for the service rules to be fully implemented. The rule action must be common for all client requests because the SR may redirect a client request to any SE in a delivery service depending on threshold conditions.

Configuring ICAP

The Internet Content Adaptation Protocol (ICAP) is an open-standards protocol that can be used for content adaptation. Content adaptation includes content translation, content filtering, and content insertion. ICAP specifies how the SE, acting as an HTTP proxy server, can communicate with an external server, which filters and adapts the requested content. This allows you to set up interoperability with Camiant PCMM-compliant third-party policy servers to allocate guaranteed bandwidth for authorized requests of content.

ICAP provides two content-processing modes for HTTP services. These modes define the transactions that can occur between an SE acting as an ICAP client and the external ICAP server. The two modes are as follows:

- Request modification (reqmod)—Allows modification of requests as they are sent from the SE to the ICAP server on their way to the origin server. The ICAP server can modify these requests depending on the content requested.
- Response modification (respmod)—Allows modification of requests after they return from the origin server. The ICAP server only acts on requested objects only after they return from the origin server.

An ICAP service is a collection of attributes that define the service and one or more ICAP servers that provide the ICAP services. You can configure a maximum of ten ICAP services per Content Engine, with an upper limit of five ICAP servers per ICAP service. Also, you can choose to apply ICAP services on all HTTP requests processed by the Content Engine or apply ICAP processing only to requests that match the Rules Template.

Caution The maximum file size supported is 2 GB. Files that exceed this size limit are not supported for ICAP processing.

To configure ICAP settings, do the following:

Step 1 Choose **Devices > Devices > Service Control > ICAP**. The ICAP page is displayed (Figure 4-12).

Figure 4-12 ICAP Page

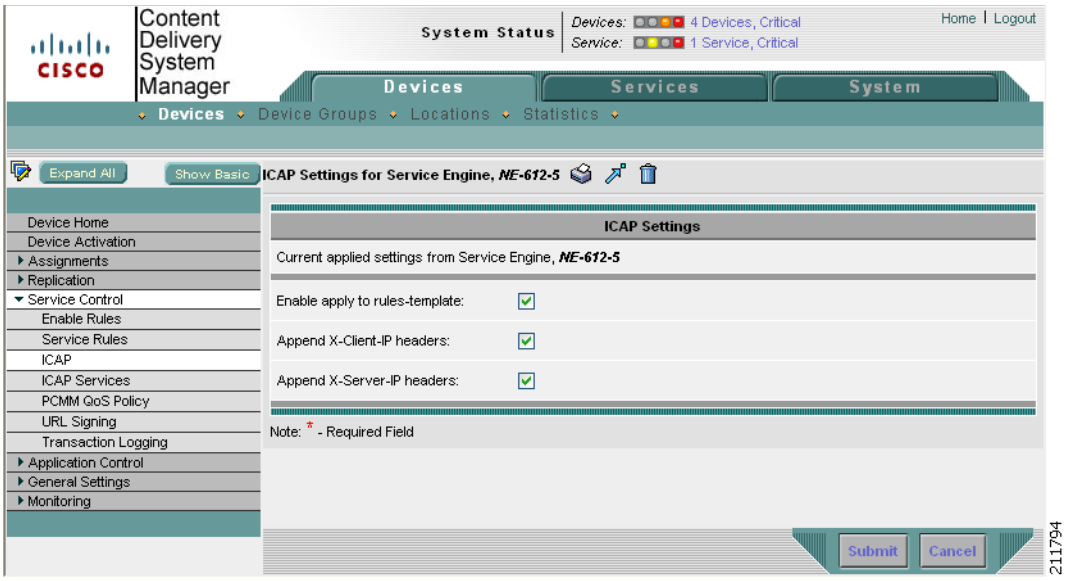

Step 2 Check the **Enable apply to rules-template** check box, if you want to apply ICAP processing only to the requests that match the use-icap-service rule from the Service Rules page.

Step 3 Check the **Append X-Client-IP headers** check box to add an X-client-IP header for ICAP processing of HTTP requests.

Step 4 Check the **Append X-Server-IP headers** check box to add an X-server-IP header for ICAP processing.

Step 5 Click **Submit** to save the settings.

Step 6 Choose **Devices > Devices > Service Control > ICAP Services**. The ICAP Services Table is displayed.

Step 7 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 8 Click the **Create New** icon in the task bar. The ICAP Services page is displayed (Figure 4-13).

To edit the ICAP service, click the **Edit** icon next to the ICAP service you want to edit.

Figure 4-13 ICAP Services Page

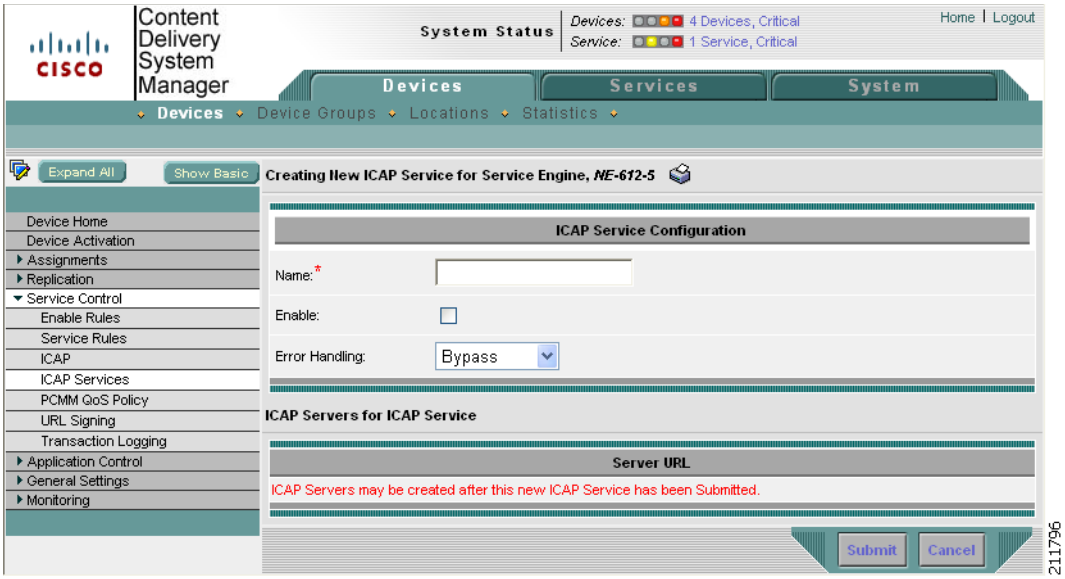

Step 9 In the **Name** field, enter a name for the ICAP service.

Step 10 Check the **Enable** check box to enable ICAP service.

Step 11 From the **Error Handling** drop-down list, choose the type of error-handling mechanism for ICAP processing. To bypass this ICAP service, choose **Bypass**. Otherwise, choose **Return Error** if you want errors to be returned for client requests. These errors are also entered in the transaction log to show the status of the action performed by the ICAP services.

Step 12 Click **Submit** to save the settings and to continue to configure an ICAP server.

Step 13 From the ICAP Services Table page, click the **Edit** icon next to the ICAP server you want to edit.

Step 14 Click the **Create New** icon in the ICAP servers area of the page. The ICAP Server page is displayed (Figure 4-14).

To edit an existing ICAP server, click the **Edit** icon next to the ICAP server.

Figure 4-14 ICAP Server Page

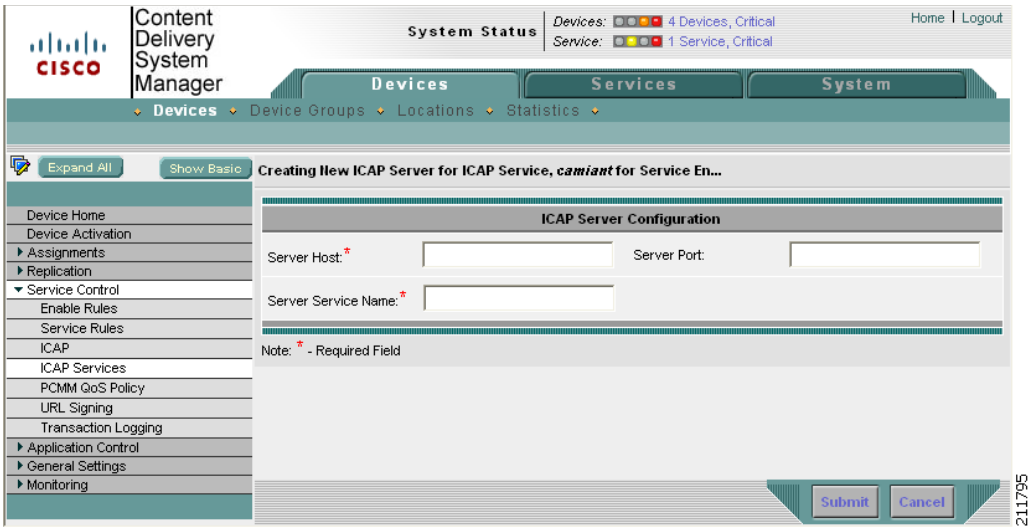

Step 15 Enter the settings as appropriate. See Table 4-8 for a description of the fields.

Table 4-8 ICAP Server Fields

Field	Description
Server Host	Hostname or IP address of the ICAP server.
Server Port	Port number on which the ICAP server is to be configured to process HTTP requests. The default port number is 1344. If no port number is specified, the default is used.
Server Service Name	Path to the ICAP server configured on the SE. Use the URL format: icap://ICAPserverIPaddress:port/service-name. The service name entered here must be supported by the ICAP vendor.

Step 16 Click **Submit** to save the settings.

Configuring PCMM QoS Policy

The SE can interoperate with Camiant PCMM-compliant third-party policy servers to allocate guaranteed bandwidth for requests from authenticated clients.

Note This is a licensed feature. Please ensure you have purchased a PCMM license for this advanced feature.

Note To use this license-enabled feature, you must have a Camiant PCMM-compliant policy server in your network and it must be configured as an ICAP server. See the "Configuring ICAP" section on page 4-19 for more information.

To configure a PCMM QoS Policy, do the following:

Step 1 Choose **Devices > Devices > Service Control > PCMM QoS Policy**. The PCMM QoS Policy page is displayed (Figure 4-15).

Figure 4-15 PCMM QoS Policy Page

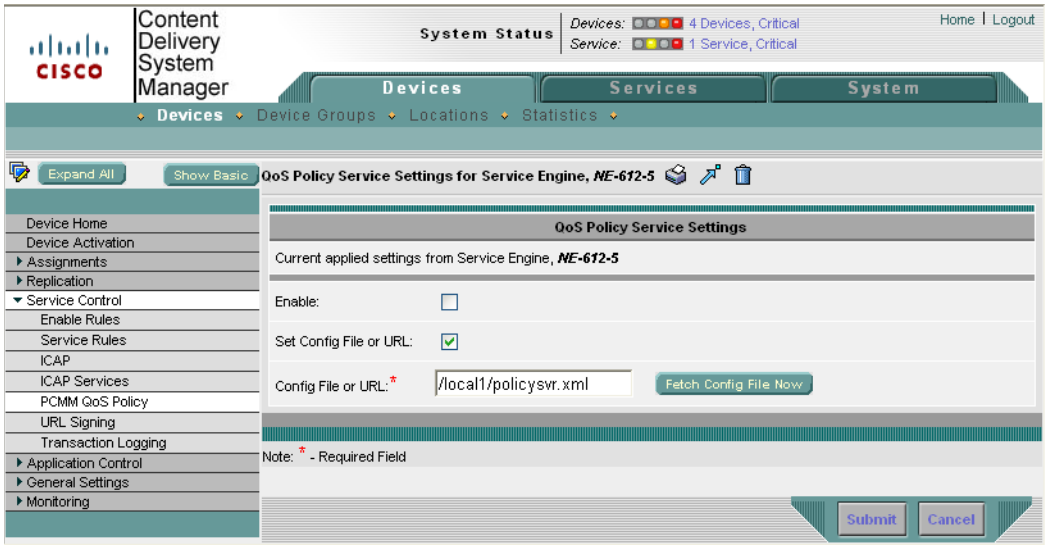

Step 2 Check the **Enable** check box to enable policy server settings for the device.

Step 3 Check the **Set Config File or URL** check box to specify the configuration file for the HTTP callout to the policy server.

The configuration file is an XML file that contains information on the callout URLs, attributes, application type, protocol, and so on.

Step 4 In the **Config File or URL** field, enter the filename or URL from which the SE can download the configuration file and click **Fetch Config File Now**. This field supports only URLs that use HTTP, HTTPS, or FTP.

Step 5 Click **Submit** to save the settings.

Configuring URL Signing

URL signature keys are word values that ensure URL-level security. The URL signature key is a shared secret between the device that assigns the key and the device that decrypts the key. Based on your network settings, either the SE itself or some other external device can assign the signature key to the URL, but the SE decrypts the URL signature key.

The CDS uses a combination of key owners, key ID numbers, and a word value to generate URL signature keys. You can have a maximum of 32 key owners. Each key owner can have up to 16 key ID numbers.

To create request-specific URL signature keys, you can choose to append the IP address of the client that has made the request to the URL signature key.

To create a URL signature key, do the following:

Step 1 Choose **Devices > Devices > Service Control > URL Signing.** The URL Signing Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click the **Create New** icon in the task bar. The URL Signing page is displayed.

To edit the URL signature, click the **Edit** icon next to the URL Signature Key ID owner you want to edit.

Step 4 Enter the settings as appropriate. See Table 4-9 for a description of the fields.

Table 4-9 URL Signature Key Settings

Field	Description
Key ID Owner	Specifies the ID number for the owner of this encryption key. Valid entries are from 1 to 32.
Key ID	Specifies the encryption key ID number. Valid entries are from 1 to 16.
URL Signature Key	Enters a unique URL signature key with up to 16 characters (excluding double quotes at the beginning and end of the string). This field does not support a space or the following special characters: pipe (|), question mark (?), double quotes ("), and apostrophe (’). Quoted and unquoted strings are allowed. Double quotes (") are allowed at the beginning and end of the string only. If you do not surround the key string with double quotes, quotes will be added when you click **Submit**.

Step 5 Click **Submit** to save the settings.

Service Rules for Directing Requests to a Policy Server

If your network is configured to work with Camiant PCMM-compliant third-party policy servers for servicing requests that require guaranteed bandwidth, you can use the following rule patterns and rule actions to filter the requests and to direct them to the policy server. The rule patterns and rule actions also enable you to generate URL signatures in the response for a valid request for a Windows Media

metafile (.asx file extension), Movie Streamer file, or Flash Media Streaming file, and to validate the URL signature on incoming requests to the SE. For more information on creating service rules, see the "Configuring Service Rules" section on page 4-15.

URL signature key authentication is implemented by using the generate-url-signature and validate-url-signature rule actions that can be applied to specific rule patterns.

Note Movie Streamer and Flash Media Streaming support URL signing in Release 2.2. Flash Media Streaming only supports the following actions: allow, block, and validate-url-signature.

Table 4-10 lists the rule patterns that support the use-icap-service rule action for directing requests that require guaranteed bandwidth to the third-party policy server:

Table 4-10 Rule Patterns Supported for use-icap-service

Rule Pattern	Description
url-regex	Filters the request based on any regular expression n the URL.
domain	Filters the request based on the domain name specified.
src-ip	Filters the request based on the IP address of the source.
header-field user-agent	Filters the request based on the user agent specified in the request header.
header-field referer	Filters the request based on the referer in the request header.
header-field request-line	Filters the request based on the request line in the request header.

You can set the se-icap-service rule action for any of the rule patterns above. If the request matches the parameters that you have set for the rule pattern, then the SE redirects the request to the third-party policy server using ICAP services. However, you must make sure that your network is configured to interoperate with the third-party policy server using ICAP services. You can set up the necessary ICAP configurations from the ICAP Services page. See the "Configuring ICAP" section on page 4-19.

You can also use the rule pattern and rule action to generate URL signatures in the response for a valid request for a Windows Media metafile. You can use the following rule patterns to filter out requests for which you want to generate a URL signature key:

Rule Pattern	Description
url-regex	Filters the request based on any regular expression in the URL.
domain	Filters the request based on the domain name specified.

For the rule patterns mentioned above, you can set the following rule actions:

Rule Action	Description
generate-url-signature	Generates the URL signatures in the Windows Media metafile response associated with pre-positioned content, based on the SE configuration for the URL signature and this rule action.
validate-url-signature	Validates the URL signature for a request by using the configuration on your SE for the URL signature and allows the request processing to proceed for this request

Note When configuring service rules, you must configure the same service rules on all SEs participating in a delivery service in order for the service rules to be fully implemented. The rule action must be common for all client requests because the SR may redirect a client request to any SE in a delivery service depending on threshold conditions.

Configuring Transaction Logs

Transaction logs allow administrators to view the traffic that has passed through the SE. Typical fields in the transaction log are the date and time when a request was made, the URL that was requested, whether it was a cache hit or a cache miss, the type of request, the number of bytes transferred, and the source IP address. For more information about transaction logs and their formats, see the "Transaction Logs" section on page 7-31.

To enable transaction logging, do the following:

Step 1 Choose **Devices > Devices > Service Control > Transaction Logs**. The Transaction Log Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-11 for a description of the fields.

Table 4-11 Transaction Log Settings Fields

Field	Description
General Settings	
Transaction Log Enable	Enables transaction logging.
Log Windows Domain	If NTLM authentication is configured, you can record the Windows domain name and username in the "authenticated username" field of the transaction log by checking this check box. For more information, see the "Transaction Logging and NTLM Authentication" section on page 7-33.
Compress Files before Export	When this check box is checked, archived log files are compressed into gzip format before being exported to external FTP servers
Log File Format Log Format Custom	Log file format choices are **extended-squid** or **apache**. For more information, see the "Transaction Log Formats" section on page 7-31. Or, choose **Log Format Custom** and enter a custom format string. For more information, see the "Custom Format" section on page 7-32.
Archive Settings	
Max size of Archive File	Maximum size (in kilobytes) of the archive file to be maintained on the local disk.
Max number of files to be archived	Maximum number of files to be maintained on the local disk.

Table 4-11 Transaction Log Settings Fields (continued)

Field	Description
Archive occurs	How often the working log is archived and the data is cleared from the working log. Choose one of the following: Choose **every** to archive every so many seconds, and enter the number of seconds for the interval. Choose **every hour** to archive using intervals of one hour or less, and choose one of the following: • **at**—Specifies the minute in which each hourly archive occurs • **every**—Specifies the number of minutes for the interval (2, 5, 10, 15, 20, or 30) Choose **every day** to archive using intervals of one day or less, and choose one of the following: • **at**—Specifies the hour in which each daily archive occurs • **every**—Specifies the number of hours for the interval (1, 2, 3, 4, 6, 8, 12, 24) Choose **every week on** to archive at intervals of one or more times a week, choose the days of the week, and choose what time each day.
Export Settings	
Enable Export	Enables exporting of the transaction log to an FTP server.
Export occurs	How often the working log is sent to the FTP server and the data is cleared from the working log. Choose one of the following: Choose **every** to export every so many seconds, and enter the number of seconds for the interval. Choose **every hour** to export using intervals of one hour or less, and choose one of the following: • **at**—Specifies the minute in which each hourly export occurs • **every**—Specifies the number of minutes for the interval (2, 5, 10, 15, 20, or 30) Choose **every day** to export using intervals of one day or less, and choose one of the following: • **at**—Specifies the hour in which each daily export occurs • **every**—Specifies the number of hours for the interval (1, 2, 3, 4, 6, 8, 12, 24) Choose **every week on** to export using intervals of one or more times a week, choose the days of the week, and what time each day.
Export Server	IP address or hostname of the FTP server.
Name	Name of the user.
Password	Password for the user.
Confirm Password	Confirms the password for the user.
Directory	Name of the directory used to store the transaction logs on the FTP server.

Table 4-11 Transaction Log Settings Fields (continued)

Field	Description
Windows Media Settings	
Enable Windows Media Settings	Enables Windows Media transaction logging.

Log File Format	Sets Windows Media Streaming Engine to generate transaction logs in the following formats:
	• **extended wms-41** Uses the standard Windows Media Services 4.1 format to generate the transaction log and includes the following three additional fields in the transaction log:
	  • SE_action (cache hit or cache miss)
	  • SE-bytes (number of bytes sent from the SE for a cache hit)
	  • username (username of the Windows Media request when NTLM, Negotiate, Digest, or basic authentication is used)
	• **extended wms-90** Uses the standard Windows Media Services 9 format to generate the transaction log and includes the following three additional fields in the transaction log:
	  • SE_action (cache hit or cache miss)
	  • SE-bytes (number of bytes sent from the Service Engine for a cache hit)
	  • username (username of the Windows Media request when NTLM, Negotiate, Digest, or basic authentication is used)
	• **wms-41** Standard Windows Media Services 4.1 format.
	• **wms-90** Standard Windows Media Services 9 format.
	For more information, see the "Windows Media Transaction Logging" section on page 7-35.

Step 3 Click **Submit** to save the settings.

Application Control

The Application Control pages provide settings for bandwidth management of delivery services and protocol engines. Configuring application control consists of the following procedures:

- Configuring Default and Maximum Bandwidth
- Configuring Bandwidth Schedules
- Configuring Windows Media Streaming—General Settings

- Configuring Windows Media Streaming Multicast Station
- Configuring Windows Media Streaming Multicast Station Schedule
- Configuring Windows Media Streaming—Broadcast Alias
- Configuring Windows Media Streaming—Bypass List
- Configuring Movie Streamer—General Settings
- Configuring RTSP Advanced Settings
- Configuring Flash Media Streaming
- Configuring Web Engine HTTP Connections
- Configuring Web Engine HTTP Caching
- Configuring Web Engine HTTP Cache Freshness
- Configuring Web Engine Advanced HTTP Caching

Configuring Default and Maximum Bandwidth

The bandwidth used for delivering content is determined by the settings in the Default and Maximum Bandwidth page, and the Scheduled Bandwidth page. The default settings are used unless a scheduled bandwidth is configured for a specified time period.

To configure the default and maximum bandwidth settings, do the following:

Step 1 Choose **Devices > Devices > Application Control > Default and Maximum Bandwidth**. The Default and Maximum Bandwidth page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-12 for a description of the fields.

Table 4-12 Application Control Default and Maximum Bandwidth Fields

Field		Description
Windows Media Incoming	Default Bandwidth	Default bandwidth allowed for incoming Windows Media traffic from client devices.
	Maximum Bandwidth	Maximum bandwidth permitted by system license. The maximum bandwidth for concurrent Windows Media streams enforces the aggregate bandwidth of all concurrent Windows Media streaming sessions, which includes RTSP-using-UDP, RTSP-using-TCP, MMS-over-HTTP, and live stream splitting. **Note** The default value, without this performance-based license, is 500 megabits per second (Mbps). The maximum allowed is 4 gigabits per second (Gbps) on a CDE200.

Table 4-12 Application Control Default and Maximum Bandwidth Fields (continued)

Field		Description
Windows Media Outgoing	Default Bandwidth	Default bandwidth allowed for outgoing Windows Media traffic from the SE.
	Maximum Bandwidth	Maximum bandwidth permitted by system license. The maximum bandwidth for concurrent Windows Media streams enforces the aggregate bandwidth of all concurrent Windows Media streaming sessions, which includes RTSP-using-UDP, RTSP-using-TCP, MMS-over-HTTP, and live stream splitting. **Note** The default value, without this performance-based license, is 500 Mbps. The maximum allowed is 4 Gbps on a CDE200.
Movie Streamer Incoming[1]	Default Bandwidth	Default bandwidth allowed for incoming Movie Streamer traffic from client devices.
	Maximum Bandwidth	Maximum bandwidth permitted by system license. The maximum bandwidth for concurrent Movie Streamer streams enforces the aggregate bandwidth of all concurrent Movie Streamer sessions. **Note** The default value, without this performance-based license, is 2 Gbps.
Movie Streamer Outgoing[1]	Default Bandwidth	Default bandwidth allowed for outgoing Movie Streamer traffic from the SE.
	Maximum Bandwidth	Maximum bandwidth permitted by system license. The maximum bandwidth for concurrent Movie Streamer streams enforces the aggregate bandwidth of all concurrent Movie Streamer sessions. **Note** The default value, without this performance-based license, is 2 Gbps.

1. Movie Streamer Incoming and Movie Streamer Outgoing fields are Release 2.2 features. For Releases 2.0 and 2.1, Movie Streamer bandwidth has default and maximum bandwidth fields that control only outgoing streaming.

Step 3 Click **Submit** to save the settings.

Configuring Bandwidth Schedules

Bandwidth Schedule settings take precedence over Default Bandwidth settings.

To configure a Bandwidth Schedule, do the following:

Step 1 Choose **Devices > Devices > Application Control > Bandwidth Schedules**. The Application Control Bandwidth Schedule Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click **Create New** in the task bar. The Scheduled Bandwidth page is displayed.

To edit a bandwidth schedule, click the **Edit** icon next to the scheduled bandwidth you want to edit.

Step 4 Enter the settings as appropriate. See Table 4-13 for a description of the fields.

Table 4-13 Application Control Bandwidth Schedule Fields

Field	Description
Bandwidth Type	Windows Media Incoming—Incoming Windows Media streaming content requests from end users. Windows Media Outgoing—Outgoing Windows Media content from SEs. Movie Streamer Incoming—Incoming Movie Streamer content requests from end users. Movie Streamer Outgoing—Outgoing Movie Streamer content in response to RTSP requests from end users. **Note** Movie Streamer Ingoing and Movie Streamer Outgoing fields are Release 2.2 features. In Releases 2.0 and 2.1, the Movie Streamer option applies to outgoing Movie Streamer content. HTTP—Sending content in response to HTTP requests from end users. This field applies to Releases 2.0 and 2.1 only.
Bandwidth Rate	Maximum amount of bandwidth you want to allow (in kilobits per second).
Start Time	Time of day for the bandwidth rate setting to start, using a 24-hour clock in local time (hh:mm).
End Time	Time of day for the bandwidth rate setting to end (hh:mm).
Use Specific Days	Days of the week on which configured bandwidth settings apply. • Full Week—Bandwidth settings are applied to the entire week. • Sun, Mon, Tue, Wed, Thu, Fri, and Sat—Specific days of the week on which configured bandwidth settings apply.
Specific Day Range	Range of days of the week on which configured bandwidth settings apply. • Start day—Day of the week to start for allowable bandwidth. • End day—Day of the week to end for allowable bandwidth.

Step 5 Click **Submit** to save the settings.

Bandwidth Graph

To view a graphical representation of the bandwidth settings, click the **Display Graph** icon in the task bar. The Application Bandwidth graph is displayed in a new window.

The vertical axis of the graph represents the amount of bandwidth in kilobits per second (kb/s) , and the horizontal axis represents the days of the week. The units shown on the vertical axis are determined dynamically based on the bandwidth rate for a particular bandwidth type. The units shown on the horizontal axis represent 24 hours per each day of the week. Each type of bandwidth is represented by a different color. A legend at the bottom of the graph maps colors to the corresponding bandwidth type.

To view the graph by bandwidth type, detailed or composite view, or days of the week, click a view option in the text at the top of the window. Table 4-14 describes the view options.

Table 4-14 Viewing Options for Content Services Bandwidth Graph

Option	Description
Windows Media In	Displays the bandwidth settings for incoming Windows Media traffic.
Windows Media Out	Displays the bandwidth settings for outgoing Windows Media traffic.
Movie Streamer In[1]	Displays the bandwidth settings for incoming Movie Streamer traffic. **Note** In Releases 2.0 and 2.1, the Movie Streamer displays outgoing streaming content.
Movie Streamer Out[1]	Displays the bandwidth settings for outgoing Movie Streamer traffic.
HTTP	Displays the bandwidth settings for HTTP requests. This field applies to Releases 2.0 and 2.1 only.
All Servers	Displays a consolidated view of all configured bandwidth types. This is the default view and is combined with the Full Week view.
Show Detailed Bandwidth/Show Effective Bandwidth	Toggles between the two options: Show Detailed Bandwidth—Displays detailed bandwidth settings for the SE and its associated device groups. The bandwidth settings of the device and device groups are shown in different colors for easy identification. Show Effective Bandwidth—Displays the composite (aggregate) bandwidth settings for the SE and its associated device groups.
Show Aggregate View/Show Non-Aggregate View	Toggles between the two options: Show Aggregate View—Displays the bandwidth settings configured for the corresponding device groups. Show Non-Aggregate View—Displays the bandwidth settings configured for the SE.
Sun, Mon, Tues, Wed, Thurs, Fri, Sat	Displays the bandwidth settings for the corresponding day of the week.
Full Week	Displays the bandwidth settings for the entire week. This is the default view and is combined with the All Servers view.

1. Movie Streamer In and Movie Streamer Out options are Release 2.2 features.

Configuring Windows Media Streaming—General Settings

To configure the General Settings for Windows Media Streaming, do the following:

Step 1 Choose **Devices > Devices > Application Control > Windows Media Streaming > General Settings**. The Windows Media Streaming General Settings page is displayed (Figure 4-16).

Figure 4-16 Windows Media Streaming Page—General Settings

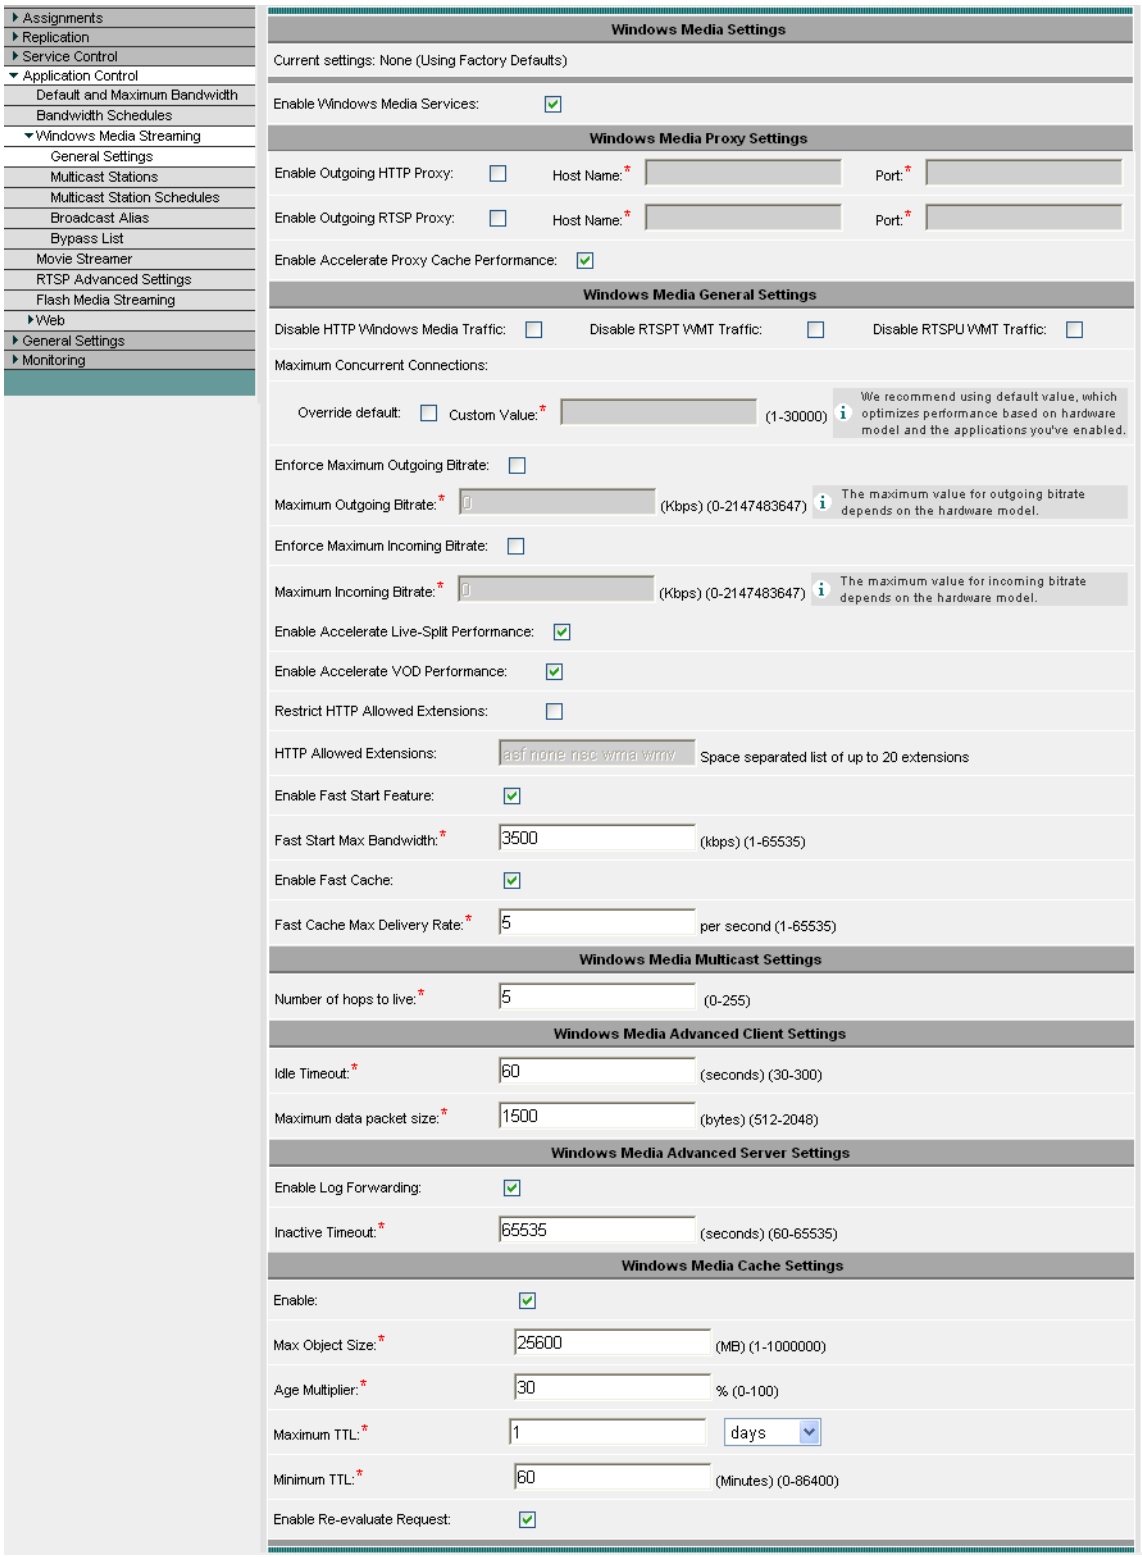

Step 2 Enter the settings as appropriate. See Table 4-15 for a description of the fields.

Table 4-15 Windows Media Streaming General Settings Fields

Field	Description
Enable Windows Media Services	When checked, Windows Media Services is enabled. To disable services, uncheck the check box.
Windows Media Proxy Settings	
Enable Outgoing HTTP Proxy	When enabled, allows an outgoing HTTP proxy server for streaming media in MMS format (MMS-over-HTTP).
Outgoing HTTP Proxy Host Name and Port	Hostname, or IP address, and port of the outgoing HTTP proxy. Valid port numbers range from 1 to 65535.
Enable Outgoing RTSP Proxy	When enabled, allows an outgoing RTSP proxy server for streaming media using RTSP.
Outgoing RTSP Proxy Host Name and Port	Hostname, or IP address, and port of the outgoing RTSP proxy. Valid port numbers range from 1 to 65535.
Enable Accelerate Proxy Cache Performance	When enabled, caching performance improvements are applied to the Windows Media proxy.
Windows Media General Settings	
Disable HTTP Windows Media Traffic	To disallow streaming over HTTP, check the check box.
Disable RTSPT WMT Traffic	To disallow streaming over RTSPT (RTSP using TCP), check the check box.
Disable RTSPU WMT Traffic	To disallow streaming over RTSPU (RTSP using UDP), check the check box.
Maximum Concurrent Connections: Override Default and Custom Value	To override the default maximum number of concurrent sessions, check the check box and enter a value in the **Custom Value** field. The default is 14000 sessions. The range is from 1 to 14000.
Enforce Maximum Outgoing Bitrate	Enforces the maximum stream bit rate for serving content when checked.
Maximum Outgoing Bitrate	The maximum streaming bit rate that can be served in kilobits per second (kbps).
Enforce Maximum Incoming Bitrate	Enforces the maximum incoming bit rate for receiving content when checked.
Maximum Incoming Bitrate	The maximum streaming bit rate (kbps) that can be received.
Enable Accelerate Live-Split Performance	Enables performance improvements in live splitting for the Windows Media proxy.
Enable Accelerate VOD Performance	Enables performance improvements in Video On Demand for the Windows Media proxy.
Restrict HTTP Allowed Extensions	Allows you to add or remove permitted extensions.

Table 4-15 Windows Media Streaming General Settings Fields (continued)

Field	Description
HTTP Allowed Extensions	List of allowable extensions for HTTP. You can add or delete filename extensions from this list with the following restrictions: • Each extension must be alphanumeric, with the first character in the extension being an alphabetic character. • You cannot have more than 10 characters in a filename extension. • You cannot add more than 6filename extensions to the allowed list.
Enable Fast Start Feature	Enables Fast Start for MMS-over-HTTP or RTSP.
Fast Start Max Bandwidth	Maximum bandwidth (kbps) allowed per Windows Media Player when Fast Start is used to serve packets to this player. The default is 3500. The range is from 1 to 65535.
Enable Fast Cache	Enables Fast Cache for MMS-over-HTTP or RTSP.
Fast Cache Max Delivery Rate	Maximum delivery rate (kbps) allowed per Windows Media Player when Fast Cache is used to deliver packets to this player. The default is 5. The range is from 1 to 65535.
Windows Media Multicast Settings	
Number of hops to live	Number of hops to live for multicast Windows Media packets. The default is 5. The range is from 0 to 255.
Windows Media Advanced Client Settings	
Idle Timeout	Number of seconds to timeout when the client connection is idle. The default is 60 The range is from 30 to 300.
Maximum Data Packet Size	Maximum packet size (in bytes) allowed. The default is 1500. The range is from 512 to 2048.
Windows Media Advanced Server Settings	
Enable Log Forwarding	Enables log forwarding to an upstream SE or Windows Media server.
Inactive Timeout	Number of seconds to timeout when the upstream SE or Windows Media server connection is idle. The default is 65535. The range is from 60 to 65535.
Windows Media Cache Settings	
Enable	When checked, Windows Media cache settings are enabled.
Max Object Size	The maximum content object size (in megabytes) the SE will cache. The default is 25600. The range is from 1 to 10000000.
Age Multiplier	The age multiplier value (as a percentage) enables the SE to estimate the life of an object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the SE. The default value is 30. The range is from 0 to 100.

Table 4-15 Windows Media Streaming General Settings Fields (continued)

Field	Description
Maximum TTL	The maximum time-to-live for objects in the cache. The value ranges are the following: 1 to 157680000 seconds 1 to 2628000 minutes 1 to 43800 hours 1 to 1825 days The default is 1 day.
Minimum TTL	The minimum time-to-live (in minutes) for objects in the cache. The default is 60. The range is from 0 to 86400.
Enable Re-evaluate Request	When checked, the cache is validated with the origin server instead of validating the cache using heuristics.

Step 3 Click **Submit** to save the settings.

Configuring Windows Media Streaming Multicast Station

This multicast-out delivery feature enables you to distribute streaming media by allowing different devices on the IP multicast to receive a single stream of content from the SE simultaneously. A more efficient method is to use the live program feature, which organizes the SEs into a delivery distribution tree, with a Content Acquirer at the root. For more information, see the "Configuring Programs" section on page 5-26.

To enable Windows Media Streaming for multicast-out, do the following:

Step 1 Choose **Devices > Devices > Application Control > Windows Media Streaming > Multicast Station**. The Multicast Station Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click **Create New** in the task bar. The Multicast Station page is displayed.

To edit a multicast station, click the **Edit** icon next to the station you want to edit.

Step 4 Enter the settings as appropriate. See Table 4-16 for a description of the fields.

Table 4-16 Multicast Station Fields

Field	Description
Station Name	Name for the multicast station.
Address	Class D IP address to be used as the multicast station IP address.
Port	The port number to be used by this station.
Media	The source URL of the Windows Media streaming media. The source file, for example *source.asf*, can be stored on any Windows Media server or SE. For multicast-in multicast-out, the source file is an .nsc file. For unicast-in unicast-out, the source file can be any Windows Media file type.
Repeat Forever	If checked, the media file plays continuously without interruption. Otherwise, the multicast stream stops when the end of the source .nsc file is reached.
Unicast URL	The URL to allow unicast live splitting for clients who cannot be reached by multicast. The unicast published URL can be used for unicast live splitting to clients that cannot receive multicast. The Windows Media Player falls back to unicast in the event of a multicast failure. The unicast published URL is made available inside the multicast description metafile (.nsc) to viewers for automatic fallback to unicast live streaming.
NSC Reference URL	If you want to use a server-side playlist as the media source for a multicast program, enter the URL for the .nsc file in this field. Make sure the .nsc file is saved in a location where it can be accessed using HTTP. Use the following format for the NSC reference URL: http://hostname/path/filename.nsc
Schedule Start Now	When checked, any active multicast stations are automatically restarted after a device reloads (reboots).
Retry Count	Number of times the multicast station retries a multicast stream from the configured source.
Retry Interval	Period of time between retries.
Alternate Source	Source to use if the primary source of the multicast fails. You can configure up to eight alternate sources for the multicast source. Enter URLs of alternates.
Enable Windows Media Multicast Logging	When enabled, the SE collects statistics on multicast streams using the .nsc file. Receives feedback on certain statistical data, such as the number of times that buffering occurred while the stream was played, the number of packets lost during transmission, the browser type used when the Windows Media Player is embedded in a browser, and the protocol used to access the stream.
Enable Local Multicast Logging	When the stream stops playing, the Windows Media Player automatically collects and sends the statistics to the multicast logging URL using the HTTP POST request method. You can choose to specify the URL to which Windows media transaction log files must be sent.
Multicast Logging URL	Either enable local multicast logging to store the log files locally, or enter a URL where to store the files.

Step 5 Click **Submit** to save the settings.

Configuring Windows Media Streaming Multicast Station Schedule

To configure the Windows Media Streaming Multicast Station Schedule, do the following:

Step 1 Choose **Devices > Devices > Application Control > Windows Media Streaming > Multicast Station Schedules**. The Windows Media Streaming Multicast Station Schedule Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click the **Create New** icon in the task bar. The Windows Media Streaming Multicast Station Schedule page is displayed.

To edit a Multicast Station Schedule, click the **Edit** icon next to the station you want to edit.

Step 4 From the **Station Name** drop-down list, choose a multicast station.

Step 5 In the **Month** field, enter the month (**1**–**12**) during which you want the schedule to start.

Step 6 In the **Day** field, enter the day (**1**–**31**) on which you want the schedule to start.

Step 7 In the **Hour** field, enter hour (**0**–**23**) that you want the schedule to start.

Step 8 In the **Minute** field, enter the minute (**0**–**59**) that you want the schedule to start.

Step 9 Click **Submit** to save the settings.

Configuring Windows Media Streaming—Broadcast Alias

The broadcast alias feature enables you to deliver an incoming stream live to requesting clients using multicast or unicast as the source of the streaming media. In this scenario, a unicast-out publishing point is created to deliver the incoming stream live to requesting clients. A more efficient method is to use the live program feature, which organizes the SEs into a delivery distribution tree, with a Content Acquirer at the root. For more information, see the "Configuring Programs" section on page 5-26.

To enable Windows Media Streaming for unicast-out by using a broadcast alias, do the following:

Step 1 Choose **Devices > Devices > Application Control > Windows Media Streaming > Broadcast Alias**. The Broadcast Alias Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click the **Create New** icon in the task bar. The Broadcast Alias page is displayed.

To edit a Broadcast Alias, click the **Edit** icon next to the name you want to edit.

Step 4 In the **Broadcast Alias** field, enter an alias for the source URL or station name to be used for the broadcast.

When a broadcast alias is configured, a client makes a request to the SE as if it were the Windows Media Server, and the SE checks to see whether the incoming stream is present. If it is, then the SE joins the stream and splits it to the new client. If the request is the first client request for this stream, the SE sends the request out to the server and then serves it to the new client.

Step 5 In the **Source URL** field, enter the URL of the station to be used for the broadcast; for example:

http://172.16.30.21/station.nsc

In this example, the multicast source URL contains the file station.nsc. A source URL in the format http://server/file.nsc tells the SE that the input source is sending out Windows Media multicast streams.

The following is an example of a unicast-in:

rtsp://wms.company.com/cotv

where wms.company.com is the name of the Windows Media Server, and cotv is the name used when the broadcast alias is created. The RTSP protocol is used to retrieve the stream.

Step 6 Click **Submit** to save the settings.

Configuring Windows Media Streaming—Bypass List

Incoming bandwidth refers to the bandwidth between a local SE and the origin server. When the SE is configured for Windows Media proxy services, incoming bandwidth usage for Video On Demand (VOD) content is unpredictable. This unpredictability is because the consumption of incoming bandwidth for VOD content can be triggered arbitrarily by an end user requesting the content. If the VOD content is not found in the SE cache, a cache miss occurs, and the Windows Media proxy must fetch the content from the origin server. The SE administrator cannot predict the incoming bandwidth usage for such events, so a large number of cache-miss VOD requests can consume all of the incoming bandwidth.

The Windows Media incoming bandwidth bypass configuration allows the administrator to configure a list of hosts that bypass the incoming bandwidth limitation. Content from a source that is listed as a host in this configuration is allowed to bypass the incoming bandwidth check for broadcast alias or multicast station content. This feature is particularly useful when the administrator wants to configure a broadcast alias or multicast station for a mission-critical live event.

To configure the list of hosts for bypassing incoming bandwidth limits, follow these steps:

Step 1 Choose **Devices > Devices > Application Control > Windows Media Streaming > Bypass List**. The Bypass List page is displayed.

Step 2 In the **Windows Media BW Incoming Bypass List** field, enter up to four IP addresses or hostnames of hosts you want to bypass the incoming bandwidth check. Separate each entry with a space.

Step 3 Click **Submit** to save the settings.

Configuring Movie Streamer—General Settings

The Movie Streamer is an open-source, standards-based, streaming server that delivers hinted MPEG-4, hinted 3GPP, and hinted MOV files to clients over the Internet and mobile networks using the industry-standard RTP and RTSP.

Note In Release 2.0, Movie Streamer was in a demonstration state for live streaming, prefetched, cached and dynamically cached content.
In Releases 2.1 to 2.3, live streaming is in full production. Prefetched, cached, and dynamically cached content remain in a demonstration state.
For details of live streaming performance for Movie Streamer, please refer to the release-specific performance bulletin.

Tip In Releases 2.0 and 2.1, to enable the Movie Streamer, navigate to the Movie Streamer page (**Devices > Devices > Application Control > Movie Streamer**), check the **Enable** check box, and click **Submit**.

In Release 2.2, to configure the general settings for Movie Streamer, do the following:

Step 1 Choose **Devices > Devices > Application Control > Movie Streamer > General Settings**. The Movie Streamer General Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-17 for a description of the fields.

Table 4-17 Movie Streamer General Settings Fields

Field	Description
Enable Movie Streamer Services	When checked, Movie Streamer Services is enabled. To disable services, uncheck the check box.
Movie Streamer Proxy Settings	
Host Name	Hostname or IP address of the proxy server for Movie Streamer.
Port	Port of the proxy server for Movie Streamer. Valid port numbers range from 1 to 65535. The default is 554.
Movie Streamer General Settings	
Maximum Concurrent Connections: Override Default and Custom Value	To override the default maximum number of concurrent sessions, check the check box and enter a value in the **Custom Value** field. The default is 8000 sessions. The range is from 1 to 16000.

Table 4-17 Movie Streamer General Settings Fields (continued)

Field	Description
Enforce Maximum Outgoing Bitrate	Enforces the maximum stream bit rate for serving content when checked.
Maximum Outgoing Bitrate	The maximum streaming bit rate that can be served in kilobytes per second (Kbps). The range is from 1 to 2147483647, depending on the hardware model.
Enforce Maximum Incoming Bitrate	Enforces the maximum incoming bit rate for receiving content when checked.
Maximum Incoming Bitrate	The maximum streaming bit rate (Kbps) that can be received. The range is from 1 to 2147483647, depending on the hardware model.
Enable Accelerate VOD Performance	Enables performance improvements in Video On Demand for the Movie Streamer proxy.
Movie Streamer Advanced Client Settings	
Idle Timeout	The Idle Timeout field and the **movie-streamer advanced client idle-timeout** command (as well as the **movie-streamer advanced client rtp-timeout** command), are only intended for performance testing when using certain testing tools that do not have full support of the RTCP receiver report. Setting these timeouts to high values causes inefficient tear-down of client connections when the streaming sessions have ended. For typical deployments, it is preferable to leave these parameters set to their defaults. The default is 60.
Movie Streamer Cache Settings	
Enable	When checked, Movie Streamer caches content on the SE and the cache settings are enabled.
Age Multiplier	The age multiplier value (as a percentage) enables the SE to estimate the life of an object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the SE. The default value is 30. The range is from 0 to 100.
Maximum TTL	The maximum time-to-live for objects in the cache. The value ranges are the following: 1 to 157680000 seconds 1 to 2628000 minutes 1 to 43800 hours 1 to 1825 days The default is 1 day.
Enable Re-evaluate Request	When checked, the cache is validated with the origin server instead of validating the cache using heuristics.

Step 3 Click **Submit** to save the settings.

Configuring RTSP Advanced Settings

To configure RTSP advanced settings for the Movie Streamer, do the following:

Step 1 Choose **Devices > Devices > Application Control > RTSP Advanced Settings**. The RTSP Advanced Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-18 for a description of the fields.

Table 4-18 RTSP Advanced Settings Fields

Field	Description
Maximum Initial Setup Delay	Maximum delay allowed (in seconds) between TCP accept and the first RTSP message from the client. The default is 10 seconds.
Maximum Request Rate	Maximum number of incoming requests per second that the RTSP gateway allows. The default is 40 requests per second.

Step 3 Click **Submit** to save the settings.

Configuring Flash Media Streaming

The Flash Media Streaming engine delivers Adobe Flash applications and video files, as well as MP3 audio files using HTTP and an Adobe proprietary protocol, RTMP. For more information, see the "Flash Media Streaming Engine" section on page 1-15.

Note Flash Media Streaming is a Release 2.1feature; therefore, this menu option is not available in Release 2.0.

Note Flash Media Streaming uses port 1935 for RTMP and RTMPE streaming. In Release 2.3, Flash Media Streaming also supports RTMPT and RTMPTE over port 80.

To enable Flash Media Streaming, do the following:

Step 1 Choose **Devices > Devices > Application Control > Flash Media Streaming > General Settings**. The Flash Media Streaming General Settings page is displayed.

Step 2 Check the **Enable Flash Media Streaming** check box.

Step 3 Enter the settings as appropriate. See Table 4-19 for a description of the fields.

Table 4-19 Flash Media Streaming Fields

Field	Description
Maximum Current Sessions	The maximum concurrent sessions the Flash Media Streaming engine supports. The range is from 1 to 15000. The default is 200.
Maximum Bandwidth	Maximum bandwidth allowed for Flash Media Streaming. The range is from 1000 to 4000000 Kbps. The default is 200000.

Step 4 Click **Submit** to save the settings.

To enable Flash Media Streaming Service Monitoring, do the following:

Step 1 Choose **Devices > Devices > Application Control > Flash Media Streaming > Service Monitoring**. The Service Monitoring page is displayed.

Step 2 Check the **Enable Service Monitoring** check box.

Service Monitoring monitors the Flash Media Streaming engine memory usage. If the memory usage reaches the 1.5 GB limit for either the Flash Media Streaming core process or the Flash Media Streaming edge process, an alarm is raised and the Service Router does not redirect any new Flash Media Streaming requests to this SE.

Step 3 Click **Submit** to save the settings.

Configuring Web Engine HTTP Connections

To configure web engine HTTP connections, do the following:

Step 1 Choose **Devices > Devices > Application Control > Web > HTTP > HTTP Connections**. The HTTP Connections page is displayed (Figure 4-17).

Figure 4-17 HTTP Connections Page

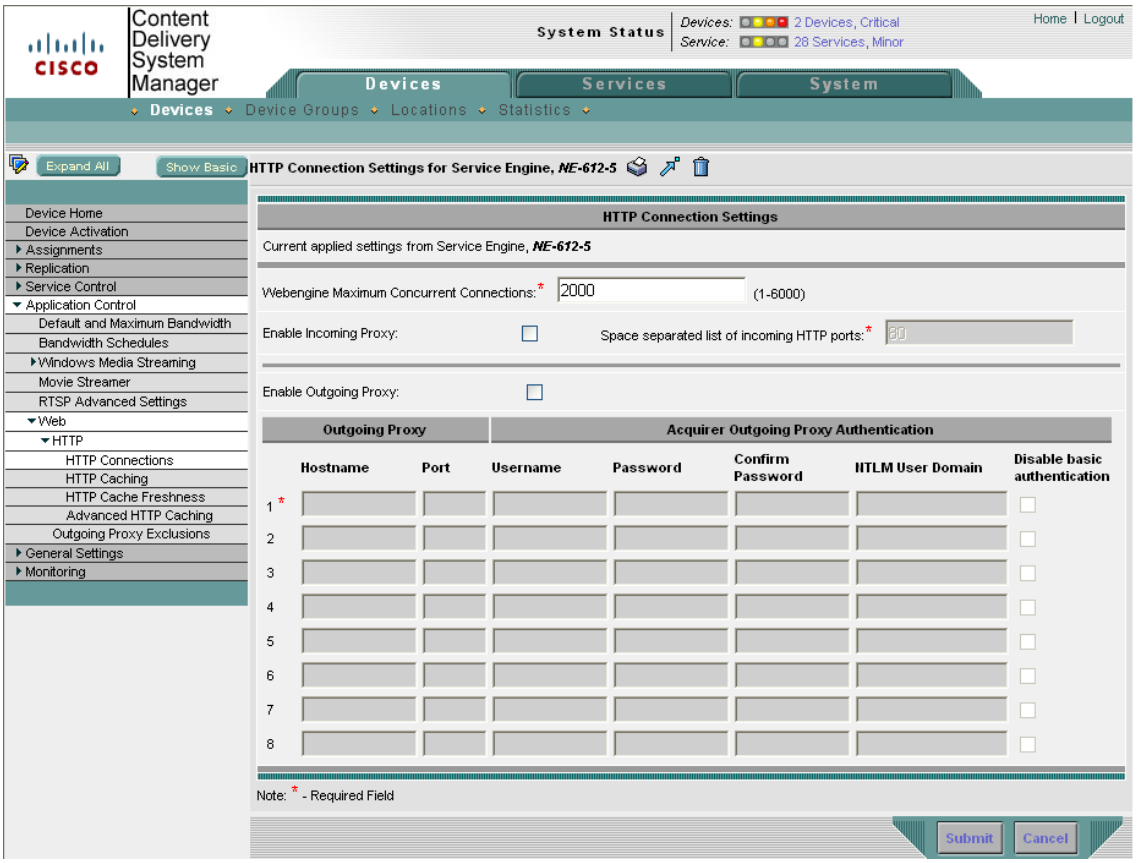

Step 2 Enter the settings as appropriate. See Table 4-20 for a description of the fields.

Table 4-20 HTTP Connections Fields

Field	Description
Webengine Maximum Concurrent Connections	The maximum concurrent sessions the Web engine supports. The range is from 100 to 6000. The default is 2000.
Enable Incoming Proxy	When enabled, accepts incoming requests on configured ports, in addition to port 80.
List of Incoming HTTP Ports	Additional ports on which to accept incoming HTTP requests. Separate port numbers with a space.
Enable Outgoing Proxy	When enabled, allows a proxy server or another SE to receive HTTP cache miss request traffic.
Content Acquirer Outgoing Proxy Authentication	
Hostname	Hostname or IP address of the outgoing proxy. The first host name or IP address entered designates that outgoing proxy server as the primary server. You can configure up to eight proxy servers. If the primary fails to respond, the request is redirected to the next proxy server.
Port	Enter the port number the proxy server uses to receive requests.

Table 4-20 HTTP Connections Fields (continued)

Field	Description
Username	If authentication is required, enter a username. This username is used for both NTLM and basic authentication.
Password	Enter the password for the user.
Confirm Password	Re-enter the password for the user.
NTLM User Domain	The NTLM server domain name used to authenticate the user.
Disable Basic Authentication	If checked, basic authentication will not occur as a fallback to an NTLM authentication failure.

Step 3 Click **Submit** to save the settings.

Configuring Web Engine HTTP Caching

To configure web engine HTTP caching, do the following:

Step 1 Choose **Devices > Devices > Application Control > Web > HTTP > HTTP Caching**. The HTTP Caching page is displayed (Figure 4-18).

Figure 4-18 HTTP Caching Page

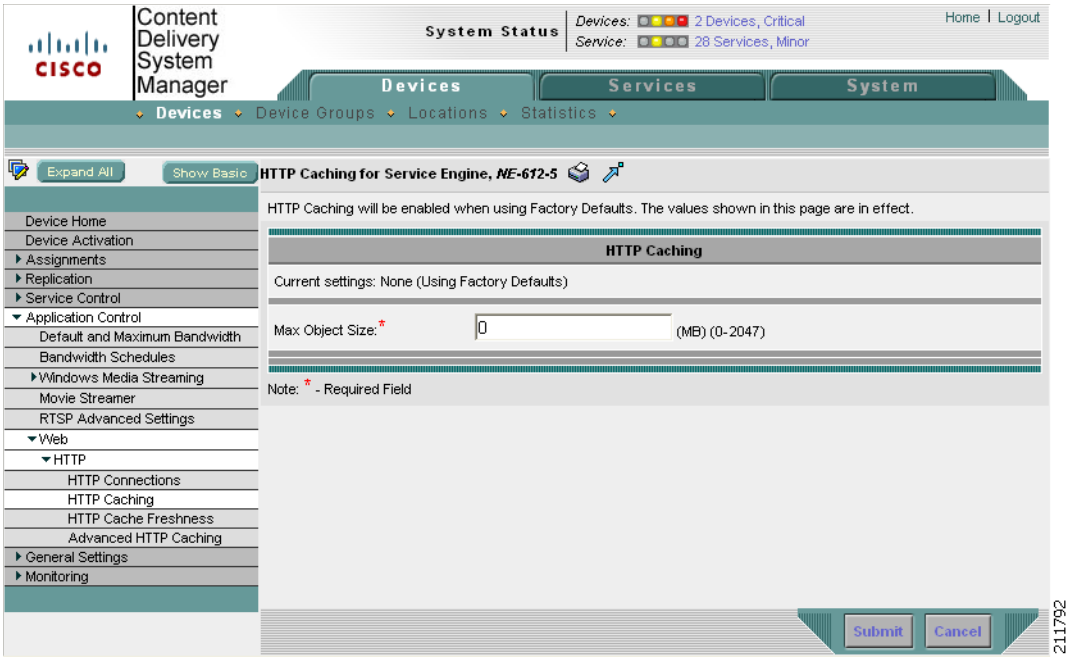

Step 2 In the Max Object Size field, enter the maximum content object size (in megabytes) that the SE will cache. The range is from 0 to 2047. Zero (0) means no limit.

Step 3 Click **Submit** to save the settings.

Configuring Web Engine HTTP Cache Freshness

To configure the web engine HTTP cache freshness, do the following:

Step 1 Choose **Devices > Devices > Application Control > Web > HTTP > HTTP Cache Freshness**. The HTTP Cache Freshness page is displayed (Figure 4-19).

Figure 4-19 HTTP Cache Freshness Page

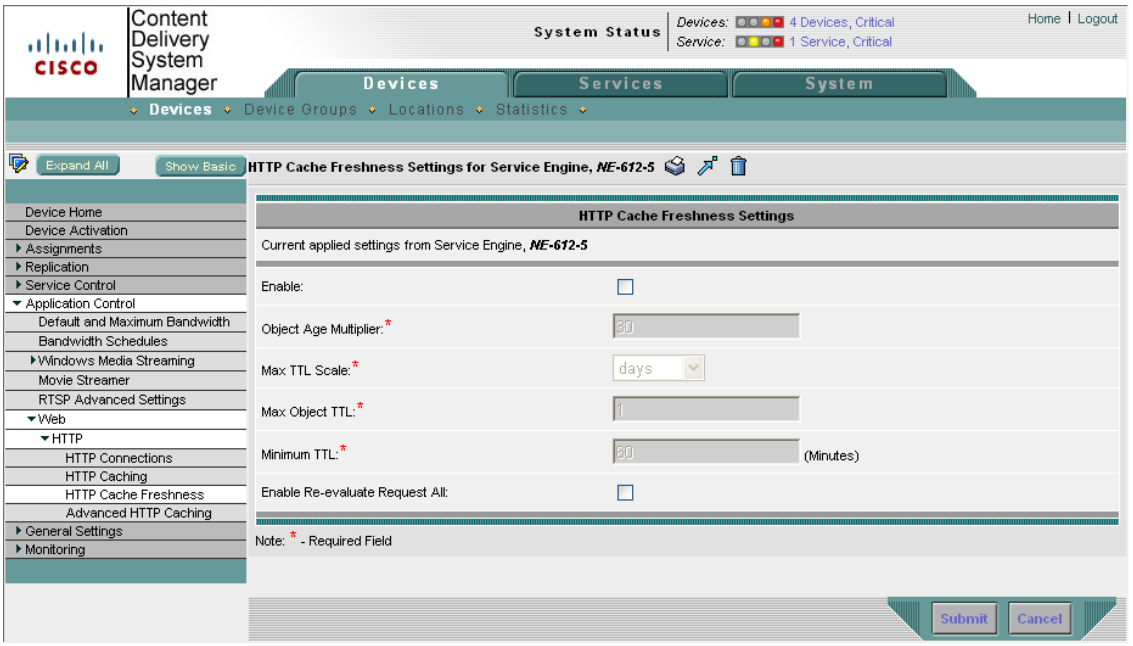

Step 2 Enter the settings as appropriate. See Table 4-21 for a description of the fields.

Table 4-21 HTTP Cache Freshness Fields

Field	Description
Enable	When checked, HTTP cache freshness is enabled.
Object Age Multiplier	The age multiplier value (as a percentage) enables the SE to guess the life of an object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the SE. The range is from 0 to 100. The default value is 30.
Max TTL Scale	The scale (seconds, hours, minutes, or days) to use for the Max Object TTL. The time-to-live (TTL) sets a ceiling on estimated expiration dates. If an object has an explicit expiration date, this takes precedence over the configured TTL.

Table 4-21 HTTP Cache Freshness Fields (continued)

Field	Description
Max Object TTL	The minimum time-to-live (TTL) for objects in cache. The valid ranges are as follows: 1 to 1825 days 1 to 43800 hours 1 to 2628000 minutes 1 to 157680000 seconds
Minimum TTL	The minimum time-to-live (in minutes) for objects in the cache. The range is from 0 to 86400. The default value is 30.
Enable Re-evaluate Request All	When enabled, all requests are re-evaluated, whether for objects or directory listings.

Step 3 Click **Submit** to save the settings.

Configuring Web Engine Advanced HTTP Caching

To configure web engine advanced HTTP cache settings, do the following:

Step 1 Choose **Devices > Devices > Application Control > Web > HTTP > Advanced HTTP Caching**. The Advanced HTTP Caching page is displayed (Figure 4-20).

Figure 4-20 Advanced HTTP Caching Page

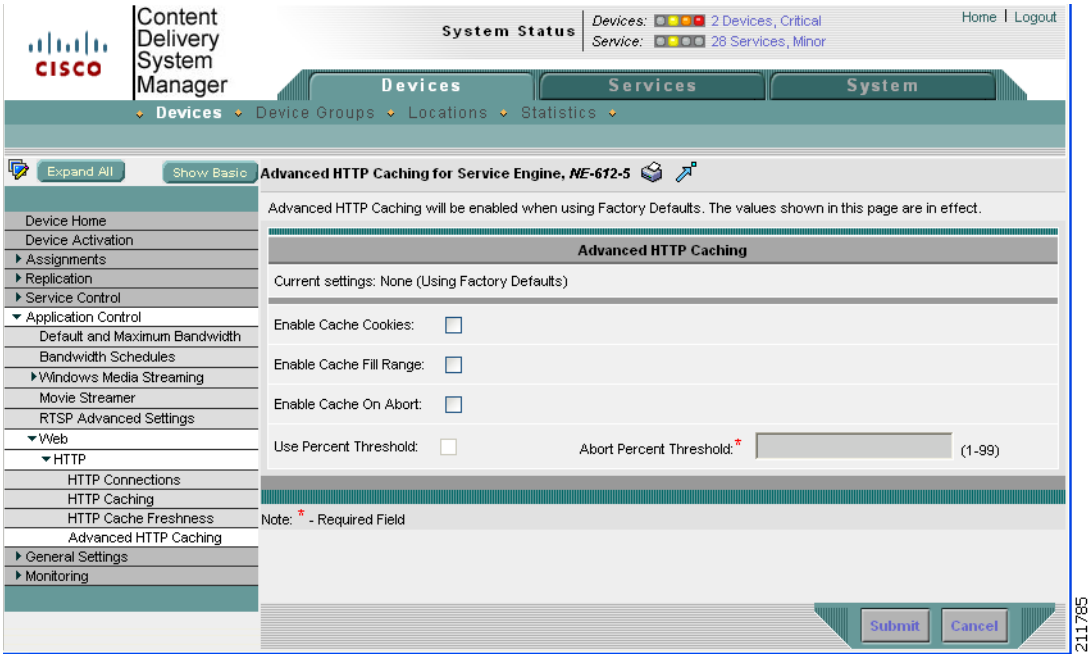

Step 2 Enter the settings as appropriate. See Table 4-22 for a description of the fields.

Table 4-22 Advanced HTTP Caching Fields

Field	Description
Enable Cache Cookies	When enabled, the Web engine caches requests with a cookie header in the response.
Enable Cache Fill Range	When enabled, the complete content is cached when the range request for the same is given by the client. The range request starts with zero (0).
Enable Cache on Abort	When enabled, caching an object is continued even though the client has aborted the request. **Note** The **Abort Percentage Threshold** field must be set for content to be cached when the client aborts.
Use Percentage Threshold	When enabled, caches an object if the percentage of the object already downloaded is greater than the percentage threshold value entered.
Abort Percentage Threshold	The percentage threshold at which the content is still cached even if the client aborts. The default value is 80 percent. **Note** To ensure content is cached when the client aborts, set the Abort Percentage Threshold to a very small value.

Step 3 Click **Submit** to save the settings.

General Settings

The General Settings pages provide settings for access control of the device, maintenance, network connectivity, and monitoring. The configuring of general settings consists of the following procedures:

- Configuring Content Management
- Login Access Control
- Authentication
- Scheduling Database Maintenance
- Setting Storage Handling
- Network Settings
- Configuring Notification and Tracking
- Configuring Troubleshooting
- Configuring Service Router Settings

Configuring Content Management

Note This is a Release 2.1 feature.

To configure the maximum number of entries for cache content, do the following:

Step 1 Choose **Devices > Devices > General Settings > Content Management**. The Content Management page is displayed.

Step 2 In the **Max Cache Content Entries** field, enter the value for the maximum entries of cached content allowed. The range is from 1 to 5000000. The default is 3000000.

Step 3 Click **Submit** to save the settings.

Login Access Control

Login authentication and authorization are used to control user access and configuration rights to CDSMs, SEs, and SRs. Login authentication is the process by which the devices verify whether the person who is attempting to log in to the device has a valid username and password. The person logging in must have a user account registered with the device. User account information serves to authorize the user for login and configuration privileges. The user account information is stored in an authentication, authorization, and accounting (AAA) database, and the devices must be configured to access the particular authentication server (or servers) where the AAA database is kept.

In a CDS network, user accounts can be created for access to the CDSM and, independently, for access to the SEs and SRs that are registered to the CDSM. For user accounts that access the CDSM, see "Configuring AAA" section on page 6-1.

Login Authentication

Login authentication provides the configuration for independent logins; in other words, login access to the device only.

Note If you plan to use a RADIUS server for authentication, you must configure RADIUS server settings before you configure and submit these settings. See the "Configuring a RADIUS Server" section on page 4-54 for RADIUS server settings information.

When the primary login server and the primary configuration server are set to local, usernames and passwords are local to each device. Local authentication and authorization uses locally configured login and passwords to authenticate login attempts.

By default, local login authentication is enabled. You can disable local login authentication only after enabling one or more of the other login authentication servers. However, when local login authentication is disabled, if you disable all other login authentication methods, local login authentication is re-enabled automatically.

Caution Make sure that RADIUS authentication is configured and operating correctly before disabling local authentication and authorization. If you disable local authentication and RADIUS is not configured correctly, or if the RADIUS server is not online, you may be unable to log in to the device.

To configure the login authentication and configuration authorization schemes for the SE, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > Login Authentication**. The Login Authentication page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-23 for a description of the fields.

Table 4-23 Login Authentication Fields

Field	Description
Enable Failover Server Unreachable	If Enable Failover Server Unreachable is enabled, the device fails over to the secondary server only if the primary server is unreachable. In any other case, for example, the authentication fails in the RADIUS server, the local database is not contacted for authentication. Conversely, if the Enable Failover Server Unreachable option is disabled, the device contacts the secondary server, regardless of the reason the authentication failed with the primary server.
Authentication Login Servers	When enabled, authentication login servers are used to authenticate user logins.
Primary Login Server	Choose either local or RADIUS.
Secondary Login Server	Choose either local or RADIUS.
Authentication Config Server	When enabled, authentication configuration servers are used to authenticate user privileges.
Primary Config Server	Choose either local or RADIUS.
Secondary Config Server	Choose either local or RADIUS.

Step 3 Click **Submit** to save the settings.

Configuring SSH

Secure Shell (SSH) consists of a server and a client program. Like Telnet, you can use the client program to remotely log in to a machine that is running the SSH server. However, unlike Telnet, messages transported between the client and the server are encrypted. The functionality of SSH includes user authentication, message encryption, and message authentication.

The SSH page allows you to specify the key length and login grace time.

To enable the SSH daemon, follow these steps:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > SSH**. The SSH page is displayed.

Step 2 Check **Enable** to enable the SSH feature. SSH enables login access to the SE through a secure and encrypted channel.

Step 3 In the **Length of Key** field, specify the number of bits needed to create an SSH key. The default is 2048.

Step 4 In the **Login Grace Time** field, specify the number of seconds the server waits for the user to successfully log in before it ends the connection. The authentication procedure must be completed within this time limit. The default is 300 seconds.

Note When changing the **Login Grace Time**, you need to first uncheck the **Enable** check box and click **Submit**. Enter the new **Login Grace Time**, check **Enable**, and click **Submit**.

Step 5 Select the SSH version.

a. To allow clients to connect using SSH protocol version 1, check the **Enable SSHv1** check box.

b. To allow clients to connect using SSH protocol version 2, check the **Enable SSHv2** check box.

Note You can enable both SSHv1 and SSHv2, or you can enable one version and not the other. You cannot disable both versions of SSH unless you disable the SSH feature by unchecking the **Enable** check box.

Step 6 Click **Submit** to save the settings.

Enabling Telnet

To enable the Telnet service, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > Telnet**. The Telnet page is displayed.

Step 2 Check **Telnet Enable** to enable the terminal emulation protocol for remote terminal connections.

Step 3 Click **Submit** to save the settings.

Setting the Message of the Day

The Message of the Day (MOTD) feature enables you to provide information bits to the users when they log in to an SE. There are three types of messages that you can set up:

- MOTD banner
- EXEC process creation banner
- Login banner

To configure the Message of the Day settings, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > Message of the Day**. The MOTD page is displayed.

Step 2 Check **Enable** to enable the MOTD settings. The Message of the Day (MOTD) banner, EXEC process creation banner, and Login banner fields become enabled.

Step 3 In the **Message of the Day (MOTD) Banner** field, enter a string that you want to display as the MOTD banner when a user attempts to log in to the SE.

Note In the Message of the Day (MOTD) Banner, EXEC Process Creation Banner, and Login Banner fields, you can enter a maximum of 980 characters. A new line character (or **Enter**) is counted as two characters, as it is interpreted as \n by the system. You cannot use special characters such as `, % ,^ , and " in the MOTD text.

Step 4 In the **EXEC Process Creation Banner** field, enter a string to be displayed as the EXEC process creation banner when a user enters into the EXEC shell of the SE.

Step 5 In the **Login Banner** field, enter a string to be displayed after the MOTD banner, when a user attempts to log in to the SE.

Step 6 Click **Submit** to save the settings.

Changing the CLI Session Time

To change the CLI session time, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > CLI Session Time**. The CLI Session Time page is displayed.

Step 2 In the CLI Session Time field, enter the time (in minutes) that the SE waits for a response before ending the session.

Step 3 Click **Submit** to save the settings.

Changing Users—Admin Password

Every device (CDSM, SE, and SR) has a built-in user account. The username is *admin* and the default password is *default*. This account allows access to all services and entities in the CDS. Any user that can access the Admin Password page in the CDSM, can configure a new password for the administrator user account on individual SEs and SRs.

To change the Admin password, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > Users > Admin Password**. The Admin Password page is displayed.

Step 2 In the **Password** field, enter a new password.

Step 3 In the **Confirm Password** field, reenter the password.

Step 4 Click **Submit** to save the settings.

Creating, Editing, and Deleting Users—Usernames

You can create, edit, and delete user accounts for login access to individual devices or device groups. A privilege profile must be assigned to each new user account. The Usernames page uses privilege profiles to determine which tasks a user can perform and the level of access provided. Users with administrative privileges can add, delete, or modify user accounts through the CDSM or the device CLI.

To create, edit, or delete a user account, do the following:

Step 1 Choose **Devices > Devices > General Settings > Login Access Control > Users > Usernames**. The User Table page is displayed.

Step 2 Aggregate Settings is set to **Yes** by default. If this SE belongs to a device group, and the settings have not been configured using the device group, you can enter the settings and they will apply to the device group as well.

If the settings are previously configured through the device group, the current configuration is displayed for view only. New settings must be configured through the associated device group.

If you want to configure settings for this SE only, and override the device group settings on this SE, set Aggregate Settings to **No**.

If you remove all device group settings, all device settings displayed with Aggregate Settings enabled are removed as well.

Note The last configuration submitted for the device, whether it is the device group configuration or the individual device configuration, is the configuration the device uses.

Step 3 Click the **Create New** icon in the task bar. The Local User page is displayed.

To edit a local user, click the **Edit** icon next to the name you want to edit.

Step 4 Enter the settings as appropriate. See Table 4-24 for a description of the fields.

Table 4-24 Local User Fields

Field	Description
Username	Name of user.
Password	User password.
Confirm Password	Re-enter user password.
Privilege	There are two types of predefined privilege profiles: • Normal user—User has read access and can see some of the SE, SR, or CDSM settings. • Superuser—User has administrative privileges such as creating new users and modifying the SE, SR, or CDSM settings.

Step 5 Click **Submit** to save the settings.

To delete a user, click the **Trash** icon in the task bar.

Authentication

User authentication and authorization (configuration rights) data can be maintained in any combination of these two databases:

- Local database (located on the SE)
- RADIUS server (external database)

The Login Authentication page allows you to choose an external access server or the internal (local) device-based authentication, authorization, and accounting (AAA) system for user access management. You can choose one method or a combination of the two methods. The default is to use the local database for authentication.

Configuring a RADIUS Server

Note The CDSM does not cache user authentication information. Therefore, the user is reauthenticated against the RADIUS server for every request. To prevent performance degradation caused by many authentication requests, install the CDSM in the same location as the RADIUS server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To configure the RADIUS server settings, do the following:

Step 1 Choose **Devices > Devices > General Settings > Authentication > RADIUS Server**. The RADIUS Server Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-25 for a description of the fields.

Table 4-25 RADIUS Server Settings Fields

Field	Description
Enable RADIUS Servers	Enables RADIUS authentication.
Time to wait	Number of seconds to wait for a response before timing out on a connection to a RADIUS server. The range is from 1 to 20. The default is 5.
Number of retransmits	Number of attempts allowed to connect to a RADIUS server. The default is 2.
Enable redirect	Redirects an authentication response to a different authentication server if an authentication request using the RADIUS server fails.
Redirect Message [1-3]	Message sent to the user if redirection occurs. **Note** If the redirect message has a space, it must be in quotes (" ").
Location [1-3]	Sets an HTML page location. This is the URL destination of the redirect message that is sent when authentication fails.
Shared Encryption Key	Encryption key shared with the RADIUS server.
Server Name [1-5]	IP address or hostname of the RADIUS server.
Server Port [1-5]	Port number on which the RADIUS server is listening.

Step 3 Click **Submit** to save the settings.

Configuring an Access Control List

To configure an access control list (ACL) for group authorization, do the following:

Step 1 Choose **Devices > Devices > General Settings > Authentication > Access Control List > Configure Access Control List**. The Access Control List Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Configure Access Control List page is displayed.

To edit a group, click the **Edit** icon next to the name you want to edit.

Step 3 Enter the settings as appropriate. See Table 4-26 for a description of the fields.

Table 4-26 Access Control List Fields

Field	Description
Action	Whether to permit or deny access for this group.
Group Name	If this action is for all groups, choose **Any Group Name**. If this action is for a specific group, choose **Enter Group Name** and enter the group name in the field.
Change Position	To change the order of this group in the access control list, which is displayed in the Access Control List Table page, click **Change Position**.

Step 4 Click **Submit** to save the settings.

To delete a group, click the **Trash** icon in the task bar.

Step 5 From the left-panel menu, choose **Enable Access Control List**. The Enable Access Control List page is displayed.

Step 6 Check the **Enable Access Control List** check box and click **Submit**.

To move a group up or down in the Access Control List table, click the Up arrow or Down arrow in the Move column.

The ACL can be applied from the SE or from a device group. The source of the currently applied settings is shown in the Access Control List Table page.

Scheduling Database Maintenance

The database maintenance runs at the scheduled time only when the following three conditions are satisfied:

- The last vacuum process happened more than 30 minutes in the past.
- The percent increase in disk space usage is greater than 10 percent.
- The available free disk space is greater than 10 percent of the total disk space.

If any of these conditions are not satisfied, the database maintenance does not run at the scheduled time.

To schedule a database cleaning or reindexing, do the following:

Step 1 Choose **Devices > Devices > General Settings > Database Maintenance**. The Database Maintenance Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-27 for a description of the fields.

Table 4-27 Database Maintenance Settings Fields

Field	Description
Full Database Maintenance Settings	
Enable	When enabled, a full database maintenance routine is performed on the SE.
Every Day	The days of the week when the maintenance is performed
Sun-Sat	When Every Day is enabled, all days of the week are also enabled.
At (time)	Time of day the maintenance is performed. Time is entered in 24-hour format as hh:mm. The default is 04:00.
Regular Database Maintenance Settings	
Enable	When enabled, a re-indexing routine is performed on the SE.
Every Day	The days of the week when the maintenance is performed.
Sun-Sat	When Every Day is enabled, all days of the week are also enabled.
At (time)	Time of day the maintenance is performed. Time is entered in 24-hour format as hh:mm. The default is 02:00.

Step 3 Click **Submit** to save the settings.

Setting Storage Handling

The Storage option offers disk error-handling settings.

Enabling Disk Error Handling

The Disk Error Handling page allows you to configure how disk errors are handled, and to define a disk device error-handling threshold.

If a bad disk drive is a critical disk drive, and the automatic reload feature (**disk error-handling reload** command) is enabled, the disk drive is marked as "bad" and the SE is automatically reloaded. After the SE is reloaded, a syslog message and an SNMP trap are generated.

The disk error-handling threshold option determines how many disk errors can be detected before the disk drive is automatically marked "bad." By default, this threshold is set to 10.

To configure a disk error-handling method, do the following:

Step 1 Choose **Devices > Devices > General Settings > Storage > Disk Error Handling**. The Disk Error Handling Settings page is displayed.

Step 2 Check the **Enable** check box.

Step 3 Check the **Enable Disk Error Handling Reload** check box if you want the device to reload the disk when a disk with a system file system (sysfs) has problems.

Step 4 Check the **Enable Disk Error Handling Remap** check box if you want to set the disks to attempt to remap disk errors automatically. This option is enabled by default.

Step 5 Check the **Enable Disk Error Handling Threshold** check box if you want to set the number of disk errors allowed before the disk is marked as bad, and enter a number (0 to 100) in the **Threshold** field. The default threshold is 10.

Step 6 Click **Submit** to save the settings.

Network Settings

The Network pages provide settings for network connectivity. Configuring network settings consist of the following procedures:

- Enabling FTP Services
- Enabling DNS
- Enabling RCP
- Configuring NTP
- Setting the Time Zone
- Viewing Network Interfaces
- Configuring External IP Addresses
- Configuring Port Channel and Load Balancing Settings
- Configuring TCP Settings
- Configuring IP General Settings
- Configuring IP ACL
- Configuring Static IP Routes

Enabling FTP Services

To enable FTP services to listen for connection requests, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > FTP**. The FTP Settings page is displayed.

Step 2 Check the **Enable FTP Services** check box.

Step 3 Click **Submit** to save the settings.

Enabling DNS

To configure Domain Name System (DNS) servers, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > DNS**. The DNS Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-28 for a description of the fields.

Table 4-28 DNS Settings Fields

Field	Description
Enable	Enables Domain Name System (DNS) on the SE.
List of DNS Servers	Space-delimited list of IP addresses for up to eight name servers for name and address resolution.
Domain Names	A space-delimited list of up to three default domain names. A default domain name allows the system to resolve any unqualified hostnames. Any IP hostname that does not contain a domain name will have the configured domain name appended to it. This appended name is resolved by the DNS server and then added to the host table. A DNS server must be configured on the system for hostname resolution to work correctly. To do this, use the List of DNS Servers field.

Step 3 Click **Submit** to save the settings.

Enabling RCP

Remote Copy Protocol (RCP) lets you download, upload, and copy configuration files between remote hosts and a switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection oriented. This service listens for requests on TCP port 514.

To enable RCP services, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > RCP**. The RCP page is displayed.

Step 2 Check the **RCP Enable** check box to have the RCP services listen for RCP requests.

Step 3 Click **Submit** to save the settings.

Configuring NTP

To configure the SE to synchronize its clock with an NTP server, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > NTP**. The NTP page is displayed.

Step 2 Check **Enable** to enable NTP.

Step 3 In the **NTP Server** field, enter the IP address or host name of up to four NTP servers. Use a space to separate the entries.

Step 4 Click **Submit** to save the settings.

Setting the Time Zone

If you have an outside source on your network that provides time services, such as an NTP server, you do not need to set the system clock manually. When manually setting the clock, enter the local time. The SE calculates Coordinated Universal Time (UTC) based on the time zone set.

Note Two clocks exist in the system: the software clock and the hardware clock. The software uses the software clock. The hardware clock is used only at startup to initialize the software clock.

Caution We highly recommend that you use NTP servers to synchronize the devices in your CDS network.

To manually configure the time zone, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > Time Zone**. The Time Zone page is displayed with the default settings of UTC (offset = 0) and no daylight savings time configured.

Step 2 To configure a standard time zone, do the following:

a. Click the **Standard Time Zone** radio button.

The standard convention for time zones uses a *Location/Area* format in which *Location* is a continent or a geographic region of the world and *Area* is a time zone region within that location. For a list of standard time zones that can be configured and their UTC offsets, see Table 4-29 on page 4-60.

b. From the **Standard Time Zone** drop-down list, choose a location for the time zone. The page refreshes, displaying all area time zones for the chosen location in the second drop-down list.

c. Choose an area for the time zone.

The UTC offset (hours and minutes ahead or behind UTC) for the corresponding time zone is displayed. During summer time savings, the offset may differ and is displayed accordingly.

Note Some of the standard time zones (mostly time zones within the United States) have daylight savings time zones configured automatically.

Step 3 To configure a customized time zone, do the following:

a. Click the **Customized Time Zone** radio button.

b. In the **Customized Time Zone** field, enter a name to for the time zone. The time zone entry is case sensitive and can contain up to 40 characters. Spaces are not allowed. If you specify any of the standard time zone names, an error message is displayed when you click **Submit**.

c. For UTC offset, choose **+** or **–** from the **UTC Offset** drop-down list to indicate whether the configured time zone is ahead or behind UTC. Also, choose the number of hours (0 to 23) and minutes (0 to 59) offset from UTC for the customized time zone. The range for the UTC offset is from –23:59 to 23:59, and the default is 0:0.

Step 4 To configure customized summer time savings, do the following:

Note Customized summer time can be specified for both standard and customized time zones.

The start and end dates for summer time can be configured in two ways: absolute dates or recurring dates. Absolute dates apply once and must be reset every year. Recurring dates apply every year.

a. Click the **Absolute Dates** radio button to configure summer settings once.

b. In the **Start Date** and **End Date** fields, specify the month, day, and year that the summer time savings starts and ends in mm/dd/yyyy format.

   Alternatively, click the **Calendar** icon and select a date. The chosen date is highlighted in blue. Click **Apply**.

c. Click the **Recurring Dates** radio button to configure a recurring summer setting.

d. Using the drop-down lists, choose the start day, week, and month when the summer time savings starts. For example, if the summer time savings begins the first Sunday in March, you would select Sunday, 1st, March from the drop-down lists.

e. Using the drop-down lists, choose the start day, week, and month when the summer time savings ends.

Step 5 Using the **Start Time** drop-down lists and the **End Time** drop-down lists, choose the hour (0 to 23) and minute (0 to 59) at which daylight savings time starts and ends.

Start Time and End Time fields for summer time are the times of the day when the clock is changed to reflect summer time. By default, both start and end times are set at 00:00.

Step 6 In the Offset field, specify the minutes offset from UTC (0 to 1439). (See Table 4-29 on page 4-60.)

The summer time offset specifies the number of minutes that the system clock moves forward at the specified start time and backward at the end time.

Step 7 To not specify a summer or daylight savings time for the corresponding time zone, click the **No Customized Summer Time Configured** radio button.

Step 8 Click **Submit** to save the settings.

Table 4-29 lists the UTC offsets for the different locations around the world.

Table 4-29 Time Zone—Offset from UTC

Time Zone	Offset from UTC (in hours)	Time Zone	Offset from UTC (in hours)
Africa/Algiers	+1	Asia/Vladivostok	+10
Africa/Cairo	+2	Asia/Yekaterinburg	+5
Africa/Casablanca	0	Asia/Yakutsk	+9
Africa/Harare	+2	Australia/Adelaide	+9.30
Africa/Johannesburg	+2	Australia/Brisbane	+10
Africa/Nairobi	+3	Australia/Darwin	+9.30
America/Buenos_Aires	–3	Australia/Hobart	+10
America/Caracas	–4	Australia/Perth	+8
America/Mexico_City	–6	Australia/Sydney	+10
America/Lima	–5	Canada/Atlantic	–4
America/Santiago	–4	Canada/Newfoundland	–3.30
Atlantic/Azores	–1	Canada/Saskatchewan	–6

Table 4-29 Time Zone—Offset from UTC (continued)

Time Zone	Offset from UTC (in hours)	Time Zone	Offset from UTC (in hours)
Atlantic/Cape_Verde	–1	Europe/Athens	+2
Asia/Almaty	+6	Europe/Berlin	+1
Asia/Baghdad	+3	Europe/Bucharest	+2
Asia/Baku	+4	Europe/Helsinki	+2
Asia/Bangkok	+7	Europe/London	0
Asia/Colombo	+6	Europe/Moscow	+3
Asia/Dacca	+6	Europe/Paris	+1
Asia/Hong_Kong	+8	Europe/Prague	+1
Asia/Irkutsk	+8	Europe/Warsaw	+1
Asia/Jerusalem	+2	Japan	+9
Asia/Kabul	+4.30	Pacific/Auckland	+12
Asia/Karachi	+5	Pacific/Fiji	+12
Asia/Katmandu	+5.45	Pacific/Guam	+10
Asia/Krasnoyarsk	+7	Pacific/Kwajalein	–12
Asia/Magadan	+11	Pacific/Samoa	–11
Asia/Muscat	+4	US/Alaska	–9
Asia/New Delhi	+5.30	US/Central	–6
Asia/Rangoon	+6.30	US/Eastern	–5
Asia/Riyadh	+3	US/East–Indiana	–5
Asia/Seoul	+9	US/Hawaii	–10
Asia/Singapore	+8	US/Mountain	–7
Asia/Taipei	+8	US/Pacific	–8
Asia/Tehran	+3.30		

The offset time (number of hours ahead or behind UTC) as displayed in the table is in effect during winter time. During summer time or daylight savings time, the offset may be different from the values in the table and is calculated and displayed accordingly by the system clock.

Viewing Network Interfaces

The Network Interfaces page is informational only. To view this information, choose **Devices > Devices > General Settings > Network > Network Interfaces**. Information about the network interfaces configured for the SE is displayed.

Configuring External IP Addresses

The External IP page allows you to configure up to eight Network Address Translation (NAT) IP addresses. This allows a router to translate up to eight internal addresses to registered unique addresses and translate external registered addresses to addresses that are unique to the private network.

To configure NAT IP addresses, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > External IP**. The External IP Settings page is displayed.

Step 2 Check the **Enable** check box.

Step 3 In the External IP Address fields (1–8), enter up to eight IP addresses.

Step 4 Click **Submit** to save the settings.

Configuring Port Channel and Load Balancing Settings

For information about configuring port channels using the CLI, see the "Redundant Dedicated Management Ports" section on page E-1.

To configure load balancing on port channels, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > Port Channel Settings**. The Port Channel Settings page is displayed.

Step 2 From the **Load Balancing Method** drop-down list, choose one of the following load balancing methods:

- **dst-ip**—Destination IP address
- **dst-mac**—Destination MAC address
- **round robin**—Each interface in the channel group

Round robin allows traffic to be distributed evenly among all interfaces in the channel group. The other balancing options give you the flexibility to choose specific interfaces (by IP address or MAC address) when sending an Ethernet frame.

Step 3 Click **Submit** to save the settings.

Configuring TCP Settings

Adjusting TCP parameters for better performance is often overlooked when configuring content caching. For data transactions and queries between client and servers, the size of windows and buffers is important, and fine-tuning the TCP stack parameters therefore becomes the key to maximizing cache performance.

Note Because of the complexities involved in TCP parameters, care is advised in tuning these parameters. In nearly all environments, the default TCP settings are adequate. Fine-tuning of TCP settings is for network administrators with adequate experience and full understanding of TCP operation details.

Caution The TCP memory limit settings allow you to control the amount of memory that can be used by the TCP subsystem send and receive buffers. The default values are device-dependent and have been chosen after extensive testing. They should not be changed under normal conditions. Increasing these values can result in the TCP subsystem using more memory, which might render the system unresponsive. Decreasing these values can result in increased response times and lower performance.

To configure TCP settings, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > TCP**. The TCP Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-30 for a description of the fields.

Table 4-30 TCP Settings Fields

Field	Description
TCP General Settings	
Enable Explicit Congestion Notification	Enables reduction of delay and packet loss. For more information, see the "Explicit Congestion Notification" section on page 4-65.
Enable Type Of Service	Enables Type of Service.
Congestion Window Size	Congestion window size in segments. The default is 2. For more information, see the "Congestion Windows" section on page 4-65.
Retransmit Time Multiplier	Factor used to modify the length of the retransmit timer by 1 to 3 times the base value determined by the TCP algorithm. For more information, see the "Retransmit Time Multiplier" section on page 4-65. **Note** Modify this factor with caution. It can improve throughput when TCP is used over slow reliable connections but should never be changed within an unreliable packet delivery environment.
Initial Slow Start Threshold	Threshold for slow start in segments. The default is 2. For more information, see the "TCP Slow Start" section on page 4-65.
Keepalive Probe Count	Number of times the SE can retry a connection before it is considered unsuccessful. The default is 4.
Keepalive Probe Interval	Length of time that the SE keeps an idle connection open. The default is 75 seconds.
Keepalive Timeout	Length of time that the SE keeps a connection open before disconnecting. The default is 90 seconds.
TCP Server Settings	
Enable Server Satellite	Sets TCP server compliance to the RFC 1323 standard. For more information, see the "TCP-Over-Satellite Extensions" section on page 4-66. **Note** If you enable a server satellite, the client satellite is automatically enabled.
Server Max Segment Size	Maximum packet size sent to the server. The default is 1460 bytes.
Server Receive Buffer Size	TCP receiving buffer size in kilobytes (1 to 512) for incoming TCP packets. The default is 32.
Server Read/Write Timeout	Period after which the SE times out when attempting to read or write to the network. The default is 120 seconds.
Server Send Buffer	TCP sending buffer size in kilobytes (1 to 512) for outgoing TCP packets. The default is 8.

Table 4-30 TCP Settings Fields (continued)

Field	Description
TCP Client Settings	
Enable Client Satellite	Sets client TCP compliance to the RFC 1323 standard. For more information, see the "TCP-Over-Satellite Extensions" section on page 4-66. **Note** If you enable a client satellite, the server satellite is automatically enabled.
Client Max Segment Size	Maximum packet size sent to clients. The default is 1432 bytes.
Client Receive Buffer Size	TCP receiving buffer size in kilobytes (8 to 512) for incoming TCP packets. The default is 8.
Client Read/Write Timeout	Period after which the SE times out when attempting to read or write to the network. The default is 120 seconds.
Client Send Buffer Size	TCP sending buffer size in kilobytes (8 to 512) for outgoing TCP packets. The default is 32.
TCP Memory Limit Settings	
TCP Memory Limit Low Water Mark	The lower limit (MB) of memory pressure mode, below which TCP enters into normal memory allocation mode. The range is from 4 to 600. Low—TCP does not regulate its memory allocation when the number of pages it has allocated globally is below this number.
TCP Memory Limit High Water Mark–Pressure	The upper limit (MB) of normal memory allocation mode, beyond which TCP enters into memory pressure mode. The range is from 5 to 610. Pressure—When the amount of memory allocated by TCP exceeds this number of pages, TCP moderates its memory consumption. This memory pressure state is exited once the number of pages allocated falls below the low mark.
TCP Memory Limit High Water Mark–Absolute	The absolute limit (MB) on TCP memory usage. The range is from 6 to 620. Absolute—The maximum number of pages, globally, that TCP will allocate.

Table 4-31 describes the default values for each command parameter, which are based on the total amount of memory for the device.

Table 4-31 Default TCP Memory Limit Settings

Total System Memory	Low	Pressure	Absolute
1 GB, 2 GB, or 4 GB	360 MB	380 MB	400 MB
512 MB	180 MB	190 MB	200 MB
256 MB	25 MB	28 MB	30 MB

Step 3 Click **Submit** to save the settings.

Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature allows an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss, including Telnet, web browsing, and transfer of audio and video data. The major issue with ECN is the need to change the operation of both the routers and the TCP software stacks to accommodate the operation of ECN.

Congestion Windows

The congestion window (*cwnd*) is a TCP state variable that limits the amount of data that a TCP sender can transmit onto the network before receiving an acknowledgment (ACK) from the receiving side of the TCP transmission. The TCP *cwnd* variable is implemented by the TCP congestion avoidance algorithm. The goal of the congestion avoidance algorithm is to continually modify the sending rate so that the sender automatically senses any increase or decrease in available network capacity during the entire data flow. When congestion occurs (manifested as packet loss), the sending rate is first lowered and then gradually increased as the sender continues to probe the network for additional capacity.

Retransmit Time Multiplier

The TCP sender uses a timer to measure the time that has elapsed between sending a data segment and receiving the corresponding ACK from the receiving side of the TCP transmission. When this retransmit timer expires, the sender (according to the RFC standards for TCP congestion control) must reduce its sending rate. However, because the sender is not reducing its sending rate in response to network congestion, the sender is not able to make any valid assumptions about the current state of the network. Therefore, in order to avoid congesting the network with an inappropriately large burst of data, the sender implements the slow start algorithm, which reduces the sending rate to one segment per transmission. (See the "TCP Slow Start" section on page 4-65.)

You can modify the sender's retransmit timer by using the Retransmit Time Multiplier field in the CDSM or the **tcp increase-xmit-timer-value** global configuration command in the CLI. The retransmit time multiplier modifies the length of the retransmit timer by one to three times the base value, as determined by the TCP algorithm that is being used for congestion control.

When making adjustments to the retransmit timer, be aware that they affect performance and efficiency. If the retransmit timer is triggered too early, the sender pushes duplicate data onto the network unnecessarily; if the timer is triggered too slowly, the sender remains idle for too long, unnecessarily slowing data flow.

TCP Slow Start

Slow start is one of four congestion-control algorithms used by TCP. The slow start algorithm controls the amount of data being inserted into the network at the beginning of a TCP session, when the capacity of the network is not known.

For example, if a TCP session began by inserting a large amount of data into the network, much of the initial burst of data would probably be lost. Instead, TCP initially transmits a modest amount of data that has a high probability of successful transmission. TCP then probes the network by sending increasing amounts of data as long as the network does not show signs of congestion.

The slow start algorithm begins by sending packets at a rate that is determined by the congestion window, or *cwnd* variable. (See the "Congestion Windows" section on page 4-65.) The algorithm continues to increase the sending rate until it reaches the limit set by the slow start threshold (*ssthresh*) variable.

(Initially, the value of the *ssthresh* variable is adjusted to the receiver's maximum window size [RMSS]. However, when congestion occurs, the *ssthresh* variable is set to half the current value of the *cwnd* variable, marking the point of the onset of network congestion for future reference.)

The starting value of the *cwnd* variable is set to that of the sender maximum segment size (SMSS), which is the size of the largest segment that the sender can transmit. The sender sends a single data segment, and because the congestion window is equal to the size of one segment, the congestion window is now full. The sender then waits for the corresponding ACK from the receiving side of the transmission. When the ACK is received, the sender increases its congestion window size by increasing the value of the *cwnd* variable by the value of one SMSS. Now the sender can transmit two segments before the congestion window is again full and the sender is once more required to wait for the corresponding ACKs for these segments. The slow start algorithm continues to increase the value of the *cwnd* variable and therefore increase the size of the congestion window by one SMSS for every ACK received. If the value of the *cwnd* variable increases beyond the value of the *ssthresh* variable, then the TCP flow control algorithm changes from the slow start algorithm to the congestion-avoidance algorithm.

TCP-Over-Satellite Extensions

The SE has the ability to turn on TCP-over-satellite extensions (as documented in RFC 1323) to maximize performance and end-to-end throughput over satellite-type connections.

The large number of satellites available to network infrastructures has increased the amount of bandwidth available in the air. Taking advantage of these connections through satellite-type connections has created new challenges in the use of TCP transactions and acknowledgments:

- Latency—Round-trip times to satellites orbiting 24,000 miles above the earth are 550 milliseconds for a single satellite hop. Buffer size must be set to prevent low-throughput connections.
- Bit errors—Packet loss can occur in a land-based device-to-satellite connection in addition to the losses caused by regular network congestion.
- Asymmetric bandwidth—Return bandwidth from satellites can be narrower than receiving bandwidth, thereby affecting performance.

Use the fields provided under the TCP Server Settings and TCP Client Settings sections to set the TCP connection so that it complies with RFC 1323.

Configuring IP General Settings

The Path Maximum Transmission Unit (MTU) Discovery discovers the largest IP packet size allowable between the various links along the forwarding path and automatically sets the correct value for the packet size. By using the largest MTU the links will bear, the sending device can minimize the number of packets it must send.

Note The Path MTU Discovery is a process initiated by the sending device. If a server does not support IP Path MTU Discovery, the receiving device has no mechanism available to avoid fragmenting datagrams generated by the server.

To enable Path MTU Discovery, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > IP General Settings**. The IP General Settings page is displayed.

Step 2 Check **Enable Path MTU Discovery**.

Step 3 Click **Submit** to save the settings.

Configuring IP ACL

Note This is a Release 2.1 feature.

Access control lists (ACLs) provide a means to filter packets by allowing a user to permit or deny IP packets from crossing specified interfaces. Packet filtering helps to control packet movement through the network. Such control can help limit network traffic and restrict network use by certain users or devices.

You can also apply ACLs to management services such as SNMP, SSH, HTTPS, Telnet, and FTP. ACLs can be used to control the traffic that these applications provide by restricting the type of traffic that the applications will handle.

In a managed CDS network environment, administrators need to be able to prevent unauthorized access to various devices and services. CDS supports standard and extended ACLs that allow administrators to restrict access to or through a CDS network device, such as the SE. Administrators can use ACLs to reduce the infiltration of hackers, worms, and viruses that can harm the network.

ACLs provide controls that allow various services to be tied to a particular interface. For example, the administrator can use IP ACLs to define a public interface on the Service Engine for content serving and a private interface for management services (for example, Telnet, SSH, SNMP, HTTPS, and software upgrades). A device attempting to access one of the services must be on a list of trusted devices before it is allowed access. The implementation of ACLs for incoming traffic on certain ports for a particular protocol type is similar to the ACL support for the Cisco Global Site Selector and Cisco routers.

To use ACLs, the system administrator must first configure ACLs and then apply them to specific services. The following are some examples of how IP ACLs can be used in various enterprise deployments:

- An application layer proxy firewall with a hardened outside interface has no ports exposed. (*Hardened* means that the interface carefully restricts which ports are available for access primarily for security reasons. Because the interface is outside, many types of attacks are possible.) The SE's outside address is globally accessible from the Internet, while its inside address is private. The inside interface has an ACL to limit Telnet, SSH, and CDSM traffic.
- An SE is deployed anywhere within the enterprise. Like routers and switches, the administrator wants to limit Telnet, SSH, and CDSM access to the IT source subnets.
- An SE is deployed as a reverse proxy in an untrusted environment, and the administrator wishes to allow only port 80 inbound traffic on the outside interface and outbound connections on the back-end interface.

Note IP ACLs are defined for individual devices only. IP ACLs cannot be managed through device groups.

When you create an IP ACL, you should note the following constraints:

- IP ACL names must be unique within the device.
- IP ACL names must be limited to 30 characters and contain no spaces or special characters.
- The CDSM can manage up to 50 IP ACLs and a total of 500 conditions per device.

- When the IP ACL name is numeric, numbers 1 through 99 denote standard IP ACLs and numbers 100 through 199 denote extended IP ACLs. IP ACL names that begin with a number cannot contain nonnumeric characters.
- Extended IP ACLs cannot be used with SNMP applications.

Creating a New IP ACL

To create a new IP ACL, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > IP ACL**. The IP ACL Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The IP ACL page is displayed.

To edit an ACL, click the **Edit** icon next to the name you want to edit.

Step 3 In the **Name** field, enter a name, observing the naming rules for IP ACLs.

Step 4 From the **ACL Type** drop-down list, choose an IP ACL type (**Standard** or **Extended**). The default is **Standard**.

Step 5 Click **Submit**. The page refreshes and the Modifying IP ACL page for a newly created IP ACL is displayed.

Note Clicking **Submit** at this point merely saves the IP ACL; IP ACLs without any conditions defined do not appear on the individual devices.

Adding Conditions to an IP ACL

To add conditions to an IP ACL, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > IP ACL**. The IP ACL Table page is displayed.

Step 2 Click the **Edit** icon next to the name of the IP ACL you want to add a condition to. The Modifying IP ACL page is displayed.

Step 3 Click the **Create New** icon in the task bar. The Condition page is displayed.

To edit a condition, click the **Edit** icon next to the name you want to edit.

Note The number of available fields for creating IP ACL conditions depends on the whether the IP ACL type is standard or extended.

Step 4 Enter values for the properties that are enabled for the type of IP ACL that you are creating.

- To create a standard IP ACL, go to Step 5.
- To create an extended IP ACL, go to Step 6.

Step 5 To set up conditions for a standard IP ACL, follow these steps:

a. From the **Purpose** drop-down list, choose a purpose (**Permit** or **Deny**).

b. In the **Source IP** field, enter the source IP address.

c. In the **Source IP Wildcard** field, enter a source IP wildcard address.

d. Click **Submit**. The Modifying IP ACL page is displayed showing the new condition and its configuration.

e. To add another condition to the IP ACL, repeat the steps.

f. To reorder your list of conditions in the Modifying IP ACL page, use the Up arrow or Down arrow in the **Order** column, or click a column heading to sort by any configured parameter.

Note The order of the conditions listed becomes the order in which IP ACLs are applied to the device.

g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL page to commit the IP ACL to the device database.

A green "Change submitted" indicator appears in the lower right corner of the Modifying IP ACL page to indicate that the IP ACL is being submitted to the device database.

Table 4-32 describes the fields in a standard IP ACL.

Table 4-32 Standard IP ACL Conditions

Field	Default Value	Description
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).
Source IP[1]	0.0.0.0	IP address of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP[1] Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

1. Required field.

Step 6 To set up conditions for an extended IP ACL, follow these steps:

a. From the **Purpose** drop-down list, choose a purpose (**Permit** or **Deny**).

b. From the **Extended Type** drop-down list, choose **Generic**, **TCP**, **UDP**, or **ICMP**.

After you choose a type of extended IP ACL, various options become available depending on what type you choose.

c. Enter the settings as appropriate. See Table 4-33 for descriptions of the extended IP ACL fields.

d. Click **Submit**. The Modifying IP ACL page is displayed showing the new condition and its configuration.

e. To add another condition to the IP ACL, repeat the steps.

f. To reorder your list of conditions from the Modifying IP ACL page, use the Up arrow or Down arrow in the **Order** column, or click a column heading to sort by any configured parameter.

Note The order of the conditions listed becomes the order in which IP ACLs are applied to the device.

g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL page to commit the IP ACL to the device database.

A green "Change submitted" indicator appears in the lower right corner of the Modifying IP ACL page to indicate that the IP ACL is being submitted to the device database.

Table 4-33 Extended IP ACL Conditions

Field	Default Value	Description	Extended Type
Purpose[1]	Permit	Specifies whether a packet is to be passed (**Permit**) or dropped (**Deny**).	Generic, TCP, UDP, ICMP
Protocol	ip	Internet protocol (**gre**, **icmp**, **ip**, **tcp**, or **udp**). To match any Internet protocol, use the **ip** keyword.	Generic
Established	Unchecked (false)	When checked, a match with the ACL condition occurs if the TCP datagram has the ACK or RST bits set, indicating an established connection. Initial TCP datagrams used to form a connection are not matched.	TCP
Source IP[1]	0.0.0.0	IP address of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.	Generic, TCP, UDP, ICMP
Source IP Wildcard[1]	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.	Generic, TCP, UDP, ICMP
Source Port 1	0	Decimal number or name of a port. Valid port numbers are 0 to 65535. See Table 4-34 and Table 4-35 for port name descriptions and associated port numbers. Valid TCP port names are as follows: • domain • exec • ftp • ftp-data • https • nfs • rtsp • ssh • telnet • www Valid UDP port names are as follows: • bootpc • bootps • domain • netbios-dgm • netbios-ns • netbios-ss • nfs • ntp • snmp • snmptrap	TCP, UDP
Source Operator	range	Specifies how to compare the source ports against incoming packets. Choices are **<**, **>**, **==**, **!=**, or **range**.	TCP, UDP
Source Port 2	65535	Decimal number or name of a port. See Source Port 1.	TCP, UDP

Table 4-33 Extended IP ACL Conditions (continued)

Field	Default Value	Description	Extended Type
Destination IP	0.0.0.0	IP address of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.	Generic, TCP, UDP, ICMP
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.	Generic, TCP, UDP, ICMP
Destination Port 1	0	Decimal number or name of a port. Valid port numbers are 0 to 65535. See Table 4-34 and Table 4-35 for port name descriptions and associated port numbers. Valid TCP port names are as follows: • domain • exec • ftp • ftp-data • https • nfs • rtsp • ssh • telnet • www Valid UDP port names are as follows: • bootpc • bootps • domain • netbios-dgm • netbios-ns • netbios-ss • nfs • ntp • snmp • snmptrap	TCP, UDP
Destination Operator	range	Specifies how to compare the destination ports against incoming packets. Choices are **<**, **>**, **==**, **!=**, or **range**.	TCP, UDP
Destination Port 2	65535	Decimal number or name of a port. See Destination Port 1.	TCP, UDP
ICMP Param Type[1]	None	Choices are **None**, **Type/Code**, or **Msg**. • **None**—Disables the ICMP Type, Code, and Message fields. • **Type/Code**—Allows ICMP messages to be filtered by ICMP message type and code. Also enables the ability to set an ICMP message code number. • **Msg**—Allows a combination of type and code to be specified using a keyword. Activates the ICMP Message drop-down list. Disables the ICMP Type field.	ICMP
ICMP Message[1]	administratively-prohibited	Allows a combination of ICMP type and code to be specified using a keyword chosen from the drop-down list. See Table 4-36 for descriptions of the ICMP messages.	ICMP
ICMP Type[1]	0	Number from 0 to 255. This field is enabled when you choose **Type/Code**.	ICMP

Table 4-33 Extended IP ACL Conditions (continued)

Field	Default Value	Description	Extended Type
Use ICMP Code[1]	Unchecked	When checked, enables the ICMP Code field.	ICMP
ICMP Code[1]	0	Number from 0 to 255. Message code option that allows ICMP messages of a particular type to be further filtered by an ICMP message code.	ICMP

1. Required field.

Table 4-34 lists the UDP keywords that you can use with extended access control lists.

Table 4-34 UDP Keywords and Port Numbers

Port Name	Description	UDP Port Number
bootpc	Bootstrap Protocol (BOOTP) client service	68
bootps	Bootstrap Protocol (BOOTP) server service	67
domain	Domain Name System (DNS) service	53
netbios-dgm	NetBIOS datagram service	138
netbios-ns	NetBIOS name resolution service	137
netbios-ss	NetBIOS session service	139
nfs	Network File System service	2049
ntp	Network Time Protocol settings	123
snmp	Simple Network Management Protocol service	161
snmptrap	SNMP traps	162

Table 4-35 lists the TCP keywords that you can use with extended access control lists.

Table 4-35 TCP Keywords and Port Numbers

Port Name	Description	TCP Port Number
domain	Domain Name System service	53
exec	Remote process execution	512
ftp	File Transfer Protocol service	21
ftp-data	FTP data connections (used infrequently)	20
https	Secure HTTP service	443
nfs	Network File System service applications	2049
rtsp	Real-Time Streaming Protocol applications	554
ssh	Secure Shell login	22
telnet	Remote login using Telnet	23
www	World Wide Web (HTTP) service	80

Table 4-36 lists the keywords that you can use to match specific ICMP message types and codes.

Table 4-36 Keywords for ICMP Message Type and Code

Message	Description
administratively-prohibited	Messages that are administratively prohibited from being allowed access.
alternate-address	Messages that specify alternate IP addresses.
conversion-error	Messages that denote a datagram conversion error.
dod-host-prohibited	Messages that signify a Department of Defense (DoD) protocol Internet host denial.
dod-net-prohibited	Messages that specify a DoD protocol network denial.
echo	Messages that are used to send echo packets to test basic network connectivity.
echo-reply	Messages that are used to send echo reply packets.
general-parameter-problem	Messages that report general parameter problems.
host-isolated	Messages that indicate that the host is isolated.
host-precedence-unreachable	Messages that have been received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 3 (Host Unreachable). This is the most common response. Large numbers of this datagram type on the network are indicative of network difficulties or may be indicative of hostile actions.
host-redirect	Messages that specify redirection to a host.
host-tos-redirect	Messages that specify redirection to a host for type of service-based (ToS) routing.
host-tos-unreachable	Messages that denote that the host is unreachable for ToS-based routing.
host-unknown	Messages that specify that the host or source is unknown.
host-unreachable	Messages that specify that the host is unreachable.
information-reply	Messages that contain domain name replies.
information-request	Messages that contain domain name requests.
mask-reply	Messages that contain subnet mask replies.
mask-request	Messages that contain subnet mask requests.
mobile-redirect	Messages that specify redirection to a mobile host.
net-redirect	Messages that are used for redirection to a different network.
net-tos-redirect	Messages that are used for redirection to a different network for ToS-based routing.
net-tos-unreachable	Messages that specify that the network is unreachable for the ToS-based routing.
net-unreachable	Messages that specify that the network is unreachable.
network-unknown	Messages that denote that the network is unknown.
no-room-for-option	Messages that specify the requirement of a parameter, but that no room is available for it.

Table 4-36 Keywords for ICMP Message Type and Code (continued)

Message	Description
option-missing	Messages that specify the requirement of a parameter, but that parameter is not available.
packet-too-big	Messages that specify that the ICMP packet requires fragmentation but the Do Not Fragment (DF) bit is set.
parameter-problem	Messages that signify parameter-related problems.
port-unreachable	Messages that specify that the port is unreachable.
precedence-unreachable	Messages that specify that host precedence is not available.
protocol-unreachable	Messages that specify that the protocol is unreachable.
reassembly-timeout	Messages that specify a timeout during reassembling of packets.
redirect	Messages that have been received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 5 (Redirect). ICMP redirect messages are used by routers to notify the hosts on the data link that a better route is available for a particular destination.
router-advertisement	Messages that contain ICMP router discovery messages called router advertisements.
router-solicitation	Messages that are multicast to ask for immediate updates on neighboring router interface states.
source-quench	Messages that have been received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 4 (Source Quench). This datagram may be used in network management to provide congestion control. A source quench packet will be issued when a router is beginning to lose packets due to the transmission rate of a source. The source quench is a request to the source to reduce the rate of a datagram transmission.
source-route-failed	Messages that specify the failure of a source route.
time-exceeded	Messages that specify information about all instances when specified times were exceeded.
timestamp-reply	Messages that contain time-stamp replies.
timestamp-request	Messages that contain time-stamp requests.
traceroute	Messages that specify the entire route to a network host from the source.
ttl-exceeded	Messages that specify that ICMP packets have exceeded the Time-To-Live configuration.
unreachable	Messages that are sent when packets are denied by an access control list; these packets are not dropped in the hardware but generate the ICMP-unreachable message.

To delete the condition, click the **Trash** icon in the task bar.

To reorder your list of conditions, use the Up arrow or Down arrow in the **Order** column, and click **Submit**.

Applying an IP ACL to an Interface

The IP ACLs can be applied to a particular interface (such as management services to a private IP address) so that the SE can have one interface in a public IP address space that serves content and another interface in a private IP address space that the administrator uses for management purposes. This feature ensures that clients can access the Service Engine only in the public IP address space for serving content and not access it for management purposes. A device attempting to access one of these applications that is associated with an IP ACL must be on the list of trusted devices to be allowed access.

To apply an IP ACL to an interface from the CLI, use the following interface configuration command:

interface {**FastEthernet** | **GigabitEthernet**} *slot/port* **ip access-group** {*accesslistnumber* | *accesslistname*} {**in** | **out**}

Deleting an IP ACL

You can delete an IP ACL, including all conditions and associations with network interfaces, or you can delete only the IP ACL conditions. Deleting all conditions allows you to change the IP ACL type if you choose to do so. The IP ACL entry continues to appear in the IP ACL listing; however, it is in effect nonexistent.

To delete an IP ACL, do the following:

Step 1 Choose **Devices > Devices**. > **General Settings > Network > IP ACL**. The IP ACL Table page is displayed.

Step 2 Click the **Edit** icon next to the name of the IP ACL that you want to delete. The Modifying IP ACL page is displayed. If you created conditions for the IP ACL, you have two options for deletion:

- **Delete ACL**—This option removes the IP ACL, including all conditions and associations with network interfaces and applications.
- **Delete All Conditions**—This option removes all the conditions, while preserving the IP ACL name.

Step 3 To delete the entire IP ACL, click **Delete ACL** (the trash can icon) in the task bar. You are prompted to confirm your action. Click **OK**. The record is deleted.

Step 4 To delete only the conditions, click **Delete All Conditions** (the trash can icon) in the task bar. You are prompted to confirm your action. Click **OK**. The window refreshes, conditions are deleted, and the ACL Type field becomes available.

Configuring Static IP Routes

The Static IP Routes page allows you to configure a static route for a network or host. Any IP packet designated for the specified destination uses the configured route.

To configure a static IP route, do the following:

Step 1 Choose **Devices > Devices > General Settings > Network > IP Routes**. The IP Route Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The IP Route page is displayed.

Step 3 In the **Destination Network Address** field, enter the destination network IP address.

Step 4 In the **Netmask** field, enter the destination host netmask.

Step 5 In the **Gateway's IP Address** field, enter the IP address of the gateway interface.

Step 6 Click **Submit** to save the settings.

Configuring Notification and Tracking

The Notification and Tracking pages provide settings for alarms, thresholds, SNMP connectivity, and device monitoring. Configuring notification and tracking consists of the following procedures:

- Enabling Alarm Overload Detection
- Setting Service Monitor Thresholds
- Configuring SNMP
- Enabling System Logs

Enabling Alarm Overload Detection

The SE tracks the rate of incoming alarms from the Node Health Manager. If the rate of incoming alarms exceeds the high-water mark (HWM) threshold, the SE enters an alarm overload state. This condition occurs when multiple applications raise alarms at the same time. When an SE is in an alarm overload state, the following events occur:

- The traps for the raise alarm-overload alarm and clear alarm-overload alarm are sent. SNMP traps for subsequent alarm raise-and-clear operations are suspended.
- Traps for alarm operations that occur between the raise-alarm-overload alarm and the clear-alarm-overload alarm operations are suspended, but individual device alarm information is still collected and available using the CLI.
- The SE remains in an alarm overload state until the rate of incoming alarms decreases to less than the low-water mark (LWM).
- If the incoming alarm rate falls below the LWM, the SE comes out of the alarm overload state and begins to report the alarm counts to the SNMP servers and the CDSM.

Alarms that have been raised on an SE can be listed by using the CLI commands shown in Table 4-37. These CLI commands allow you to systematically drill down to the source of an alarm.

Table 4-37 Viewing SE Alarms

Command	Syntax	Description
show alarms		Displays a list of all currently raised alarms (critical, major, and minor alarms) on the SE.
	show alarms critical	Displays a list of only currently raised critical alarms on the SE.
	show alarms major	Displays a list of only currently raised major alarms on the SE.
	show alarms minor	Displays a list of only currently raised minor alarms on the SE .
	show alarms detail	Displays detailed information about the currently raised alarms.

Table 4-37 Viewing SE Alarms (continued)

Command	Syntax	Description
	show alarms history	Displays a history of alarms that have been raised and cleared on the SE. The CLI retains the last 100 alarm raise and clear events only.
	show alarms status	Displays the counts for the currently raised alarms on the SE. Also lists the alarm-overload state and the alarm-overload settings.

To configure the alarm overload detection, do the following:

Step 1 Choose **Devices > Devices > General Settings > Notification and Tracking > Alarm Overload Detection**. The Alarm Overload Detection page is displayed.

Step 2 Uncheck the **Enable Alarm Overload Detection** check box if you do not want to configure the SE to suspend alarm raise and clear operations when multiple applications report error conditions. Alarm overload detection is enabled by default.

Step 3 In the **Alarm Overload Low Water Mark** field, enter the number of alarms per second for the clear alarm overload threshold. The low water mark is the level to which the number of alarms must drop below before alarm traps can be sent. The default value is 1.

Step 4 In the **Alarm Overload High Water Mark** field, enter the number of alarms per second for the raise alarm-overload threshold. The high-water mark is the level the number of alarms must exceed before alarms are suspended. The default value is 10.

Step 5 Click **Submit** to save the settings.

Setting Service Monitor Thresholds

The Service Monitor page is where you configure workload thresholds for the SE. In load-based routing, these thresholds are used to determine the best SE to serve requested content. For more information about load-based routing, see the "Configuring the Service Router" section on page 4-87.

To configure workload thresholds, do the following:

Step 1 Choose **Devices > Devices > General Settings > Notification and Tracking > Service Monitor**. The Service Monitor page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-38 for a description of the fields.

Table 4-38 Service Monitor Fields

Field	Description
CPU Settings	
Enable	Allows the SR to collect CPU load information from the SE.
Threshold	Value (as a percentage) that determines when the SE is overloaded. The threshold determines the extent of CPU usage allowed. The range is from 1 to 100. The default is 80.

Table 4-38 Service Monitor Fields (continued)

Field	Description
Sample Period	Time interval (in seconds) between two consecutive samples. The sample period is the time during which the SE and the SR exchange keep-alive messages that contain the SE load information. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recently sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
Disk Settings[1]	
Enable	Allows the SR to collect disk transaction information from the SE.
Threshold	The threshold, as a percentage, determines the extent of disk usage allowed. The range is from 1 to 100. The default is 80.
Sample Period	Time interval (in seconds) between two consecutive samples. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recently sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
Memory Settings	
Enable	Allows the SR to collect memory usage information from the SE.
Threshold	The threshold (in percent) determines the extent of memory usage allowed. The range is from 1 to 100. The default is 90.
Sample Period	Time interval (in seconds) between two consecutive samples. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recently sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
KMemory Settings[1]	
Enable	Allows the SR to collect kernel memory usage information from the SE.
Threshold	The threshold determines the extent of kernel memory usage (in percent) allowed. The range is from 1 to 100. The default is 50.
Sample Period	Time interval (in seconds) between two consecutive samples. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recently sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
WMT Settings	
Enable	Allows the SR to collect stream count information from the SE.
Threshold	Percentage of streams for which the SE has been either configured or licensed. The range is from 1 to 100. The default is 90.
Sample Period	Time interval (in seconds) between two consecutive samples. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recent sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
FMS Settings[1]	
Enable	Allows the SR to collect stream count information from the SE.
Threshold	Percentage of streams for which the SE has been either configured or licensed. The range is from 1 to 100. The default is 90.

Table 4-38 Service Monitor Fields (continued)

Field	Description
Sample Period	Time interval (in seconds) between two consecutive samples. The range is from 1 to 60. The default is 5.
Number of Samples	Number of most recently sampled values used when calculating the average. The range is from 1 to 120. The default is 6.
Movie Streamer Settings[1]	
Enable	Allows the SR to collect stream count information from the SE.
Threshold	Percentage of streams for which the SE has been either configured or licensed. The range is from 1 to 100. The default is 90.
Web Engine Settings[1]	
Enable	Allows the SR to collect stream count information from the SE.
Threshold	Percentage of streams for which the SE has either configured or licensed. The range is from 1 to 100. The default is 90. **Note** We recommend setting the Web Engine Threshold to 80 percent or higher when Flash Media Streaming is enabled.
Disk Failure Count Settings[2]	
Threshold	When the number of failed disks exceeds this threshold, no further requests are sent to this SE. The Disk Failure Count Threshold is only for the CDNFS disks. The maximum threshold that can be configured is determined by the physical number of CDNFS disks (up to 15) installed on the SE. The default is 1. **Note** When an alarm is received for a SYSFS disk, it is immediately marked as a failed disk. It is not checked against the Disk Failure Count Threshold.

1. This is a Release 2.1 feature. Sample period and number of samples are not required for Movie Streamer and Web Engine, because these protocol engines do not support bandwidth-based threshold monitoring.
2. This is a Release 2.2 feature.

Step 3 Click **Submit** to save the settings.

Configuring SNMP

The Cisco CDS supports the following versions of SNMP:

- Version 1 (SNMPv1)
- Version 2 (SNMPv2)
- Version 3 (SNMPv3)

To configure the SNMP settings, do the following:

Step 1 Choose **Devices > Devices > General Settings > Notification and Tracking > SNMP > General Settings**. The SNMP General Settings page is displayed.

Step 2 Enable the settings as appropriate. See Table 4-39 for a description of the fields.

Table 4-39 SNMP General Settings Fields

Field	Description
Traps	
Enable SNMP Settings	Enables the SNMP agent to transmit traps to the SNMP server.
Service Engine	Enables SNMP traps: • Disk Read—Enables disk read error trap. • Disk Write—Enables disk write error trap. • Disk Fail—Enables disk failure error trap. • Transaction Logging—Enables transaction log write error trap.
SNMP	Enables SNMP-specific traps: • Authentication—Enables authentication trap. • Cold Start—Enables cold start trap.
SE Alarm	Enables alarm traps: • Raise Critical—Enables raise-critical alarm trap. • Clear Critical—Enables clear-critical alarm trap. • Raise Major—Enables raise-major alarm trap. • Clear Major—Enables clear-major alarm trap. • Raise Minor—Enables raise-minor alarm trap. • Clear Minor—Enables clear-minor alarm trap.
Entity	Enables SNMP entity traps.
Config	Enables CiscoConfigManEvent error traps.
Miscellaneous Settings	
Notify Inform	Enables the SNMP notify inform request.

Step 3 Click **Submit** to save the settings.

Step 4 From the left-panel menu, click **Community**. The SNMP Community Table page is displayed.

The maximum number of community strings that can be created is ten.

Step 5 Click the **Create New** icon in the task bar. The SNMP Community page is displayed.

Click the **Edit** icon next to the community name to edit a community setting.

Step 6 Enter the settings as appropriate. See Table 4-40 for a description of the fields.

Table 4-40 SNMP Community Fields

Field	Description
Community	Community string used as a password for authentication when you access the SNMP agent of the SE. The "Community Name" field of any SNMP message sent to the SE must match the community string defined here in order to be authenticated. You can enter a maximum of 64 characters in this field.
Group name/rw	Group to which the community string belongs. The **Read/Write** option allows a read or write group to be associated with this community string. The **Read/Write** option permits access to only a portion of the MIB subtree. Choose one of the following three options from the drop-down list: • **None**—Choose this option if you do not want to specify a group name to be associated with the community string. • **Read/Write**—Choose this option if you want to allow read-write access to the group associated with this community string. • **Group**—Choose this option if you want to specify a group name.
Group Name	Name of the group to which the community string belongs. You can enter a maximum of 64 characters in this field. This field is available only if you have chosen the **Group** option in the Group name/rw field.

Step 7 Click **Submit** to save the settings.

Step 8 From the left-panel menu, click **Group**. The SNMP Group Table page is displayed.

The maximum number of groups that can be created is ten.

Step 9 Click the **Create New** icon in the task bar. The SNMP Group page is displayed.

Click the **Edit** icon next to the Group Name to edit a group.

Step 10 Enter the settings as appropriate. See Table 4-41 for a description of the fields.

Table 4-41 SNMP Group Fields

Field	Description
Name	Name of the SNMP group. You can enter a maximum of 256 characters.
Sec Model	Security model for the group. Choose one of the following options from the drop-down list: • **v1**—Version 1 security model (SNMP Version 1 [noAuthNoPriv]). • **v2c**—Version 2c security model (SNMP Version 2 [noAuthNoPriv]). • **v3-auth**—User security level SNMP Version 3 (AuthNoPriv). • **v3-noauth**—User security level SNMP Version 3 (noAuthNoPriv). • **v3-priv**— User security level SNMP Version 3 (AuthPriv).
Read View	Name of the view (a maximum of 64 characters) that enables you only to view the contents of the agent. By default, no view is defined. In order to provide read access to users of the group, a view must be specified.

Table 4-41 SNMP Group Fields (continued)

Field	Description
Write View	Name of the view (a maximum of 64 characters) that enables you to enter data and configure the contents of the agent. By default, no view is defined.
Notify View	Name of the view (a maximum of 64 characters) that enables you to specify a notify, inform, or trap. By default, no view is defined.

Step 11 Click **Submit** to save the settings.

Step 12 From the left-panel menu, click **User**. The SNMP User Table page is displayed.

The maximum number of users that can be created is ten.

Step 13 Click the **Create New** icon in the task bar. The SNMP User page is displayed.

Click the **Edit** icon next to the username to edit a user.

Step 14 Enter the settings as appropriate. See Table 4-42 for a description of the fields.

Table 4-42 SNMP User Fields

Field	Description
Name	String representing the name of the user (256 characters maximum) who can access the SE.
Group	Name of the group (256 characters maximum) to which the user belongs.
Remote SNMP ID	Globally unique identifier for a remote SNMP entity. To send an SNMPv3 message to the SE, at least one user with a remote SNMP ID must be configured on the SE. The SNMP ID must be entered in octet string format.
Authentication Algorithm	Authentication algorithm that ensures the integrity of SNMP packets during transmission. Choose one of the following three options from the drop-down list: • **No-auth**—Requires no security mechanism to be turned on for SNMP packets. • **MD5**—Provides authentication based on the hash-based Message Authentication Code Message Digest 5 (HMAC-MD5) algorithm. • **SHA**—Provides authentication based on the hash-based Message Authentication Code Secure Hash (HMAC-SHA) algorithm.
Authentication Password	String (256 characters maximum) that configures the user authentication (HMAC-MD5 or HMAC-SHA) password. The number of characters is adjusted to fit the display area if it exceeds the limit for display. This field is optional if the **no-auth** option is chosen for the authentication algorithm. Otherwise, this field must contain a value.
Confirmation Password	Authentication password for confirmation. The re-entered password must be the same as the one entered in the Authentication Password field.
Private Password	String (256 characters maximum) that configures the authentication (HMAC-MD5 or HMAC-SHA) parameters to enable the SNMP agent to receive packets from the SNMP host. The number of characters is adjusted to fit the display area if it exceeds the limit for display.
Confirmation Password	Private password for confirmation. The re-entered password must be the same as the one entered in the Private Password field.

Step 15 Click **Submit** to save the settings.

Step 16 To define a SNMPv2 MIB view, click **View** from the left-panel menu. The SNMP View Table page is displayed.

The maximum number of SNMPv2 views that can be created is ten.

Step 17 Click the **Create New** icon in the task bar. The SNMP View page is displayed.

Click the **Edit** icon next to the username to edit a view.

Step 18 Enter the settings as appropriate. See Table 4-43 for a description of the fields.

Table 4-43 SNMP View Fields

Field	Description
Name	String representing the name of this family of view subtrees (256 characters maximum). The family name must be a valid MIB name such as ENTITY-MIB.
Family	Object identifier (256 characters maximum) that identifies a subtree of the MIB.
View Type	View option that determines the inclusion or exclusion of the MIB family from the view. Choose one of the following two options from the drop-down list: • **Included**—The MIB family is included in the view. • **Excluded**—The MIB family is excluded from the view.

Step 19 Click **Submit** to save the settings.

Step 20 From the left-panel menu, click **Host**. The SNMP Host Table page is displayed.

The maximum number of hosts that can be created is four.

Step 21 Click the **Create New** icon in the task bar. The SNMP Host page is displayed.

Click the **Edit** icon next to the hostname to edit a host.

Step 22 Enter the settings as appropriate. See Table 4-44 for a description of the fields.

Table 4-44 SNMP Host Fields

Field	Description
Trap Host	Hostname or IP address of the SNMP trap host that is sent in SNMP trap messages from the SE.
Community/User	Name of the SNMP community or user (256 characters maximum) that is sent in SNMP trap messages from the SE.
Authentication	Security model to use for sending notification to the recipient of an SNMP trap operation. Choose one of the following options from the drop-down list: • **No-auth**—Sends notification without any security mechanism. • **v2c**—Sends notification using Version 2c security. • **Model v3-auth**—Sends notification using SNMP Version 3 (AuthNoPriv). • **Security Level v3-noauth**—Sends notification using SNMP Version 3 (NoAuthNoPriv security). • **Level v3-priv**—Sends notification using SNMP Version 3 (AuthPriv security).

Table 4-44 SNMP Host Fields (continued)

Field	Description
Retry	Number of retries (1 to 10) allowed for the inform request. The default is 2.
Timeout	Timeout for the inform request in seconds (1 to 1000). The default is 15.

Step 23 Click **Submit** to save the settings.

Step 24 From the left-panel menu, click **Asset Tag**. The SNMP Asset Tag page is displayed.

Step 25 In the **Asset Tag Name** field, enter a name for the asset tag.

Step 26 From the left-panel menu, click **Contact**. The SNMP Contact page is displayed.

Step 27 In the **Contact** field, enter a name of the contact person for this SE.

Step 28 In the **Location** field, enter a location of the contact person for this SE.

Step 29 Click **Submit** to save the settings.

Supported MIBs

The SNMP agent supports the following MIBs:

- ENTITY-MIB (RFC 2037)
- MIB-II (RFC 1213)
- HOST-RESOURCES-MIB (RFC 2790)
- CISCO-ENTITY-ASSET-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-SERVICE-ENGINE-MIB (supports streaming media-related MIB objects)

ENTITY-MIB, MIB-II, and HOST-RESOURCES-MIB are public-available MIBs.

The CISCO-SERVICE-ENGINE-MIB is extended to incorporate MIB objects related to streaming. The WMT and Movie Streamer groups incorporate statistics about the WMT server or proxy, and Movie Streamer. The Flash Media Streaming group incorporates statistics about the Flash Media Streaming protocol engine. For each 64-bit counter MIB object, a 32-bit counter MIB object is implemented so that SNMP clients using SNMPv1 can retrieve data associated with 64-bit counter MIB objects. The MIB objects of each of these groups are read-only.

- The WMT MIB group provides statistics about WMT proxy and server performance. Twenty-eight MIB objects are implemented in this group. Six of these MIB objects are implemented as 64-bit counters.
- The Movie Streamer MIB group provides statistics about RTSP streaming engine performance. Seven MIB objects are implemented in this group. Two of these MIB objects are implemented as 64-bit counters.
- The Flash Media Streaming MIB group provides statistics about HTTP and RTMP streaming engine performance.

The CISCO-SERVICE-ENGINE-MIB is available from your Cisco account team.

Use the following link to access the CISCO-ENTITY-ASSET-MIB and the CISCO-CONFIG-MAN-MIB:

ftp://ftp.cisco.com/pub/mibs/v2/

Note If your browser is located behind a firewall or you are connecting to the Internet with a DSL modem and you are unable to access this file folder, you must change your web browser compatibility settings. In the Internet Explorer (IE) web browser, choose **Tools > Internet Options > Advanced**, and check the **Use Passive FTP** check box.

Enabling System Logs

Use the System Logs page to set specific parameters for the system log file (syslog). This file contains authentication entries, privilege level settings, and administrative details. System logging is always enabled. By default, the system log file is stored as /local1/syslog.txt.

To enable system logging, do the following:

Step 1 Choose **Devices > Devices > General Settings > Notification and Tracking > System Logs**. The System Log Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-45 for a description of the fields.

Table 4-45 System Logs Settings Fields

Field	Description
System Logs	
Enable	Enables system logs.
Facility	Facility where the system log is sent.
Console Settings	
Enable	Enable sending the system log to the console.
Priority	Severity level of the message that should be sent to the specified remote syslog host. The default priority is warning. The priorities are: Emergency—System is unusable. Alert—Immediate action needed. Critical—Critical condition. Error—Error conditions. Warning—Warning conditions. Notice—Normal but significant conditions. Information—Informational messages. Debug—Debugging messages.
Disk Settings	
Enable	Enables saving the system logs to disk.
File Name	Path and filename where the system log file is stored on the disk. The default is /local1/syslog.txt.
Priority	Severity level of the message that should be sent to the specified remote syslog host.
Recycle	The maximum size of the system log file before it is recycled. The default is 10000000 bytes.

Table 4-45 System Logs Settings Fields (continued)

Field	Description
Host Settings	
Enable	Enables sending the system log file to a host. You can configure up to four hosts.
Hostname	A hostname or IP address of a remote syslog host.
Priority	Severity level of the message that should be sent to the specified remote syslog host.
Port	The destination port on the remote host. The default is 514.
Rate Limit	The message rate per second. To limit bandwidth and other resource consumption, messages can be rate limited. If this limit is exceeded, the remote host drops the messages. There is no default rate limit, and by default all system log messages are sent to all syslog hosts.

Step 3 Click **Submit** to save the settings.

Multiple Hosts for System Logging

Each syslog host can receive different priority levels of syslog messages. Therefore, you can configure different syslog hosts with a different syslog message priority code to enable the SE to send varying levels of syslog messages to the four external syslog hosts.

However, if you want to achieve syslog host redundancy or failover to a different syslog host, you must configure multiple syslog hosts on the SE and assign the same priority code to each configured syslog host.

Configuring Troubleshooting

The Kernel Debugger troubleshooting page allows you to enable or disable access to the kernel debugger. Once enabled, the kernel debugger is automatically activated when kernel problems occur.

Enabling the Kernel Debugger

To enable the kernel debugger, do the following:

Step 1 Choose **Devices > Devices > General Settings > Troubleshooting > Kernel Debugger**. The Kernel Debugger window appears.

Step 2 To enable the kernel debugger, check the **Enable** check box, and click **Submit**.

For information about monitoring the SEs, see the "Device Monitoring" section on page 7-6.

Configuring Service Router Settings

Note This is a Release 2.1 feature.

The keep-alive interval is used by the SE to send keep-alive messages to the SR.

To configure the keep-alive interval the SE uses for messages to this SR, do the following:

Step 1 Choose **Devices > Devices > General Settings > Service Routing Settings**. The Service Routing Settings page is displayed.

Step 2 In the **Keepalive-Interval** field, enter the number of seconds the messages from the SR should be kept alive on this SE. The range is from 1 to 120. The default is 2.

Step 3 Click **Submit** to save the settings.

Configuring the Service Router

Configuring a Service Router (SR) consists of the following procedures:

- Activating a Service Router
- Configuring Routing Settings
- IP-Based Redirection
- Configuring Application Control
- General Settings—Last Resort

For information on configuring all general settings, except Last Resort, see the "General Settings" section on page 4-48.

Activating a Service Router

Activating an SR can be done through the Device home page initially, or through the Device Activation page.

To activate an SR from the Device Activation page, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed.

Step 2 Click the **Edit** icon next to the SR you want to configure. The Device home page is displayed.

Step 3 Click **Show All** to display the top-level menu options, and choose **Device Activation**. The Device Activation page is displayed.

Step 4 Enter the settings as appropriate. See Table 4-46 for a description of the fields.

Table 4-46 Service Router Activation Fields

Field	Description
Name	Name of the device.
Activate	To activate or deactivate the device, check or uncheck the **Activate** check box.
Location	The Location drop-down list lists all the location configured for the CDS.

Table 4-46 Service Router Activation Fields (continued)

Field	Description
Coverage Zone File	To have a local Coverage Zone file overwrite the CDS network-wide Coverage Zone file, choose a file from the **Coverage Zone** drop-down list. See the "Coverage Zone File Registration" section on page 6-9 for information about creating and registering a Coverage Zone file. Otherwise, choose **None**.
Use SR's primary IP address	Enables the CDSM to use the IP address on the primary interface of the SR for management communications. **Note** If the **Use SR's primary IP Address for Management Communication** check box is checked and the Management Communication Address and Port are configured, the CDSM uses the SR's primary IP address for communication. If communication attempts to the primary address fail, the CDSM tries the configured IP address and port. **Note** Do not check the **Use SR's primary IP Address for Management Communication** check box if you want to separate management and streaming traffic. Instead, use the Management Communication Address and Port fields to specify where management traffic should be sent.
Management Communication Address	Manually configures a management IP address for the CDSM to communicate with the SR. **Note** This is a Release 2.1 feature. Manual configuration of the management IP address and port are used when using port channel configuration to separate management and streaming traffic. For more information about port channel configuration see the "Configuring Port Channel and Load Balancing Settings" section on page 4-62 and the "Configuring Port Channel" section on page E-1.
Management Communication Port	Port number to enable communication between the CDSM and the SR. **Note** This is a Release 2.1 feature.
Comments	Information about the settings.

Step 5 Click **Submit** to save the settings.

Configuring Routing Settings

To configure the SR routing method, do the following:

Step 1 Choose **Devices > Devices > Routing Settings**. The Routing Settings page is displayed.

Step 2 Enter the settings as appropriate. See Table 4-47 for a description of the fields.

Table 4-47 Service Router Routing Settings Fields

Field	Description
Least Loaded Routing	**Note** This is a Release 2.0 and Release 2.1 feature. In Release 2.2 and Release 2.3, Least Loaded Routing is enabled by default. When enabled, the SR redirects the client request to the SE that reports the lowest average load, given that each SE in the routing table has the same metric value (weight). When a configured threshold for an SE is exceeded, the SR compares the SE assigned weight and current load to the configured threshold values when making routing decisions. The metric value is configured in the Coverage Zone file. For more information about the SR routing methods, see the "Routing Methods" section on page 1-22. Last-resort routing is applicable when load-based routing (least loaded) is enabled and all Service Engines have exceeded their thresholds or all Service Engines in the domain are offline. The Service Router can redirect requests to a configurable alternate domain when all Service Engines serving a client network region are overloaded. **Note** If the last-resort domain is not configured and the Service Engine thresholds are exceeded, requests are redirected to the origin server. To configure last-resort routing, see the "General Settings—Last Resort" section on page 4-91 for Release 2.1, or see the "Configuring Last-Resort Routing" section on page E-6 for Release 2.0.
Location Based Routing	In Release 2.0 and Release 2.1, location-based routing is designed to work with least-loaded routing. When both are enabled, the Service Router first looks up the client's IP address in the Coverage Zone file. Then, if there is no subnet match, the client's geographical location is compared to the geographical location of the Service Engines listed in the Coverage Zone file, and the closest and least-loaded Service Engine is selected. Geographically locating a client is used when users roam outside of their home networks.
Primary Geo-Location Server IP Address and Port	The IP address and port number of the primary Geo-Location Server for location-based routing. For more information, see the "Routing Methods" section on page 1-22.
Secondary Geo-Location Server IP Address and Port	The IP address and port number of the secondary Geo-Location Server for location-based routing. For more information, see the "Routing Methods" section on page 1-22.
Enable Content Based Routing	When enabled, the SR redirects requests based on the URI. Requests for the same URI are redirected to the same SE, provided the SE's thresholds have not been exceeded. This optimizes disk usage in the CDS by storing only one copy of the content on one SE, instead of multiple copies on several SEs. For more information about content-based routing, see the "Content-Based Routing" section on page 1-24.
Number of Redundant Copies	Number of copies of a content to keep among SEs in a delivery service. The range is from 1 to 4. The default is 1. If redundancy is configured with more than one copy, multiple Service Engines are picked for a request with the same URI hash.

Step 3 Click **Submit** to save the settings.

IP-Based Redirection

Note IP-based redirection is a Release 2.3 feature.

IP-based redirection uses IP addresses to route client requests to the SR and on to the SE. For more information, see the "IP-Based Redirection" section on page 1-20.

To enable IP-based redirection, do the following:

Step 1 Choose **Devices > Devices > IP-based Redirection**. The IP-based Redirection page is displayed.

Step 2 Check the **Enable IP-based Redirection** check box and click **Submit**.

Configuring Application Control

The Application Control pages allow you to enable Flash Media Streaming and to enable HTTP proxy on an SR.

Note The Application Control pages for the SR are part of Release 2.1.

To configure the application control for the SR, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed.

Step 2 Click the **Edit** icon next to the SR you want to configure. The Device home page is displayed.

Step 3 Click **Show All** to display the top-level menu options, and choose **Application Control**.

Step 4 To enable Flash Media Streaming on the SR, click **Flash Media Streaming**. The Flash Media Settings page is displayed.

a. Check the **Enable Flash Media Streaming** check box.

b. Click **Submit**.

Step 5 To enable HTTP proxy on the SR, click **Web > HTTP > HTTP Connections**. The HTTP Connections Settings page is displayed.

a. Check the **Enable Incoming Proxy** check box.

b. Enter the port numbers that receive HTTP in the associated field.

Separate each port number by a space. The default is port 80.

c. Click **Submit** to save the settings.

General Settings—Last Resort

For information on configuring all general settings, except Last Resort, see the "General Settings" section on page 4-48.

Note This is a Release 2.1 feature. To configure last resort routing for Release 2.0, see the "Configuring Last-Resort Routing" section on page E-6.

Last-resort routing is applicable when load-based routing is enabled and all Service Engines have exceeded their thresholds or all Service Engines in the domain are offline. The Service Router can redirect requests to a configurable alternate domain when all Service Engines serving a client network region are overloaded.

Note If the last-resort domain is not configured and the Service Engine thresholds are exceeded, requests are redirected to the origin server.

To configure last-resort routing, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed.

Step 2 Click the **Edit** icon next to the SR you want to configure. The Device home page is displayed.

Step 3 Click **Show All** to display the top-level menu options, and choose **General Settings > Last Resort**. The Last Resort Table page is displayed.

Step 4 Click the **Create New** icon.

Click the **Edit** icon next to the domain name to edit a table entry.

Step 5 In the **Domain Name** field, enter the service routing domain name; for example, srfqdn.cisco.com.

Step 6 In the **Alternate Domain Name** field, enter where to route requests for the specified domain; for example, www.cisco.com.

Step 7 Click **Submit** to save the settings. The entry is added to the Last Resort Table.

Configuring the CDSM

Configuring a CDSM consists of the General Settings menu items. For information on configuring general settings, see the "General Settings" section on page 4-48.

Device activation is accomplished during installation and initialization of the CDS devices. See the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide* for more information. The Device Activation page for the CDSM displays information about the management IP address and the role of the CDSM.

For information about primary and standby CDSMs, see the "Configuring Primary and Standby CDSMs" section on page 3-8.

CHAPTER 5

Configuring Services

This chapter provides information on configuring the CDS. This chapter covers the following major topics:

- Configuring Delivery Services, page 5-1
- Configuring Programs, page 5-26
- Viewing Programs, page 5-39
- Copying a Program, page 5-41

Configuring Delivery Services

Delivery services are configured for prefetch ingest, hybrid ingest, and live programs. Dynamic ingest, the other type of ingest, is dynamically cached upon retrieving content that is not locally stored. For more information about content ingest types, see the "Ingest and Distribution" section on page 1-3.

Configuring a delivery service consists of defining the following:

- Content Origins
- Creating Delivery Service
- Identifying Content

Content Origins

Content is stored on origin servers. Each delivery service is configured with one origin server. The same origin server can be used by multiple live delivery services. However, only one prefetch/caching delivery service is allowed per origin server.

Note When creating a live delivery service with the same Content Origin as a prefetch/caching delivery service, the same set of SEs must be assigned to both; otherwise, the SR may redirect requests to unassigned SEs.

For more information about origin servers, see the "Origin Servers" section on page 2-5.

To create a Content Origin, do the following:

Step 1 Choose **Services > Service Definition > Content Origins**. The Content Origin Table page is displayed (Figure 5-1).

Figure 5-1 Content Origin Table Page

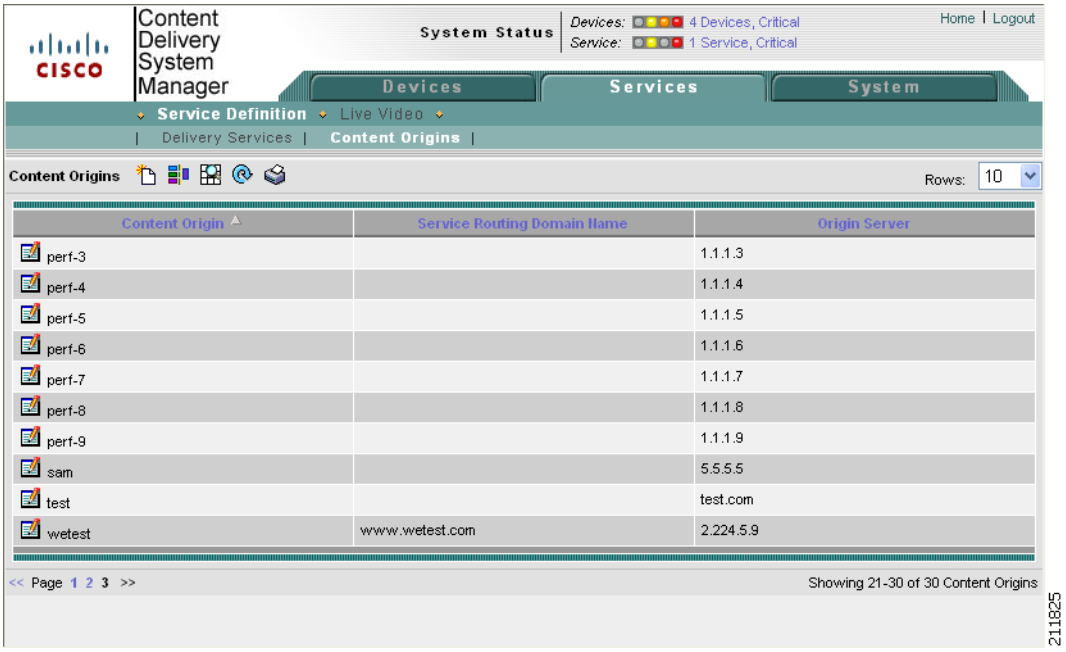

Step 2 Click the **Create New** icon in the task bar. The Content Origin page is displayed (Figure 5-2).

Figure 5-2 Content Origin Page

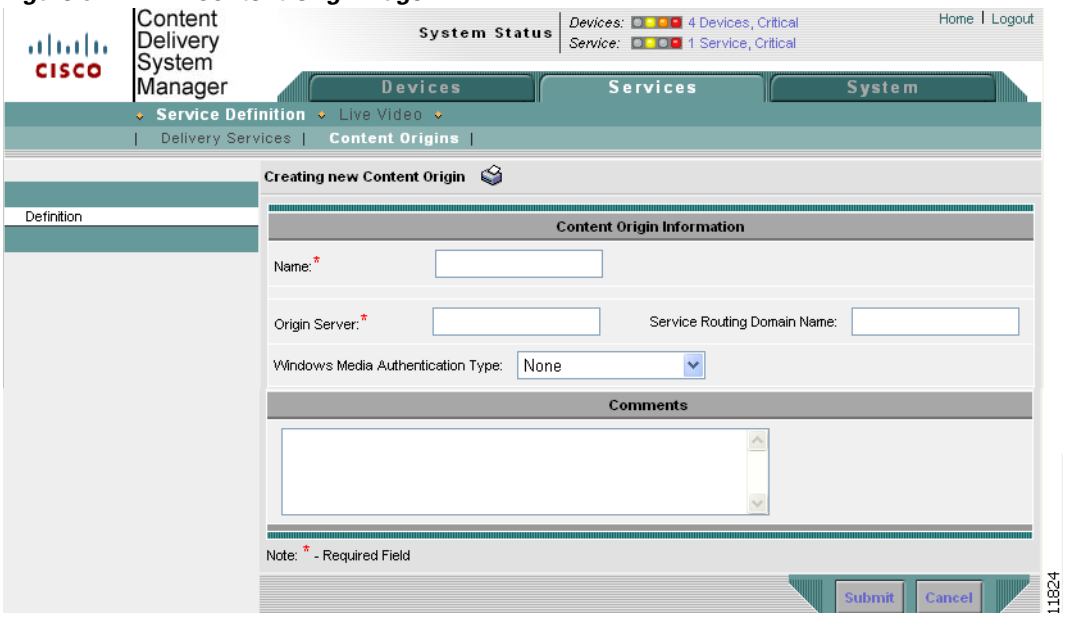

To edit a Content Origin, click the **Edit** icon next to the Content Origin name.

Step 3 Enter the settings as appropriate. See Table 5-1 for a description of the fields.

Table 5-1 Content Origin Fields

Field	Description
Name	Unique name of the origin server.
Origin Server	Fully qualified domain name (FQDN) of the origin server.
Service Routing Domain Name	The FQDN to route client requests. The SE translates the service routing domain name (SRDN) to the origin server whenever it needs to retrieve content from the origin server. The service routing domain name configured for the Content Origin should also be configured in the DNS servers, so client requests can get redirected to a Service Router for request mediation and redirection. The URLs that are published to the users have the service routing domain names as the prefix.
Windows Media Authentication Type	The type of client authentication that is required by the origin server. The options are: • None • Basic authentication • NTLM authentication • Digest\ • Negotiate
Comments	Information about the Content Origin.

Step 4 Click **Submit** to save the settings.

To delete a Content Origin, from the Content Origin Table page, click the **Edit** icon next to the Content Origin you want to delete, and click the **Trash** icon in the task bar.

Creating Delivery Service

A delivery service is a configuration used to define how content is acquired, distributed, and stored in advance of a client request. For more information about delivery services, see the "Delivery Service" section on page 2-3.

Before creating delivery services, make sure the devices that participate in the delivery service are configured for the type of content to be delivered.

A delivery service configuration consists of the following:

1. Service Definition
2. SE and Content Acquirer Assignment or Device Group and Content Acquirer Assignment
3. PCMM Configuration
4. General Settings
5. Session Shifting
6. Identifying Content

Configuring the first five are described in the following procedure. Identifying content is described in the "Identifying Content" section on page 5-11.

Tip For information about testing a delivery service, see Appendix G, "Testing the Internet Streamer CDS."

To create a delivery service, do the following:

Service Definition

Step 1 Choose **Services > Service Definition > Delivery Services**. The Delivery Services Table page is displayed (Figure 5-3).

Step 2 Click the **Create New** icon in the task bar. The Delivery Services Definition page is displayed (Figure 5-3).

To edit a delivery service, click the **Edit** icon next to the delivery service name.

Figure 5-3 Delivery Service Definition Page

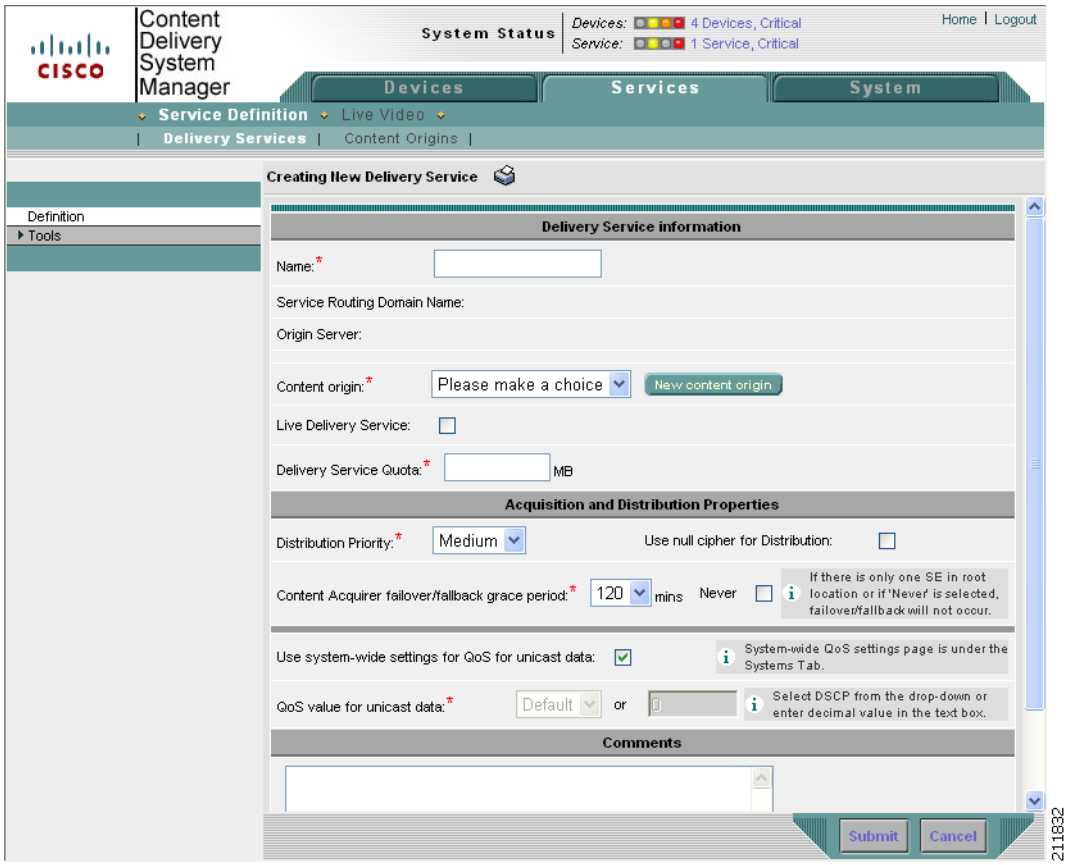

Step 3 Enter the settings as appropriate. See Table 5-2 for a description of the fields.

Table 5-2 Delivery Service Definition Fields

Field	Description
Delivery Service Information	
Name	Unique name for the delivery service.
Content Origin	All Content Origins that have been created are listed in the drop-down list. The delivery service and the Content Origin have a one-to-one relationship. To create a new Content Origin, click **New Content Origin**.
Live Delivery Service	When checked, creates a live program to distribute live or scheduled programs to the SEs associated with this delivery service and with the live program. This delivery service does not have a related Manifest file and cannot be used to distribute file-based content as regular delivery services do. The live program learns about a live stream through a program file that describes the attributes of the program. Checking this check box disables the Delivery Service Quota field and fields in the Acquisition and Distribution Properties section.
Delivery Service Quota	Maximum content storage size, in megabytes, for prefetched, hybrid, or, cached content for this delivery service.

Table 5-2 Delivery Service Definition Fields (continued)

Field	Description
Acquisition and Distribution Properties	
Distribution Priority	Content distribution priority setting. Options are High, Normal, and Low. The default is Normal. The priority of content acquisition also depends on the origin server. Requests from different origin servers are processed in parallel. Requests from the same origin server are processed sequentially by their overall priority.
Use null cipher for Distribution	When checked, disables encryption for distribution.
Content Acquirer failover/fallback grace period	Number of minutes before a Content Acquirer failover or a temporary Content Acquirer fallback occurs. The range is from 20 to 120 minutes. For more information, see the "Content Acquirer Redundancy" section on page 1-27.
Never	When checked, SE failover or fallback never occurs.
Use system-wide settings for QoS for unicast data	When checked, applies the system-wide QoS settings for unicast data to the delivery service. To override the system-wide QoS settings with delivery service-specific QoS values, leave this check box unchecked, and configure the delivery service-specific QoS values in the QoS value for unicast data field.
QoS value for unicast data	Configures a Differentiated Services Code Point (DSCP) value for the QoS. If you choose **Other**, enter a decimal value in the corresponding field. You can set QoS settings on a per-delivery service basis and a system-wide global configuration basis. Delivery service settings take precedence over global settings.
QoS value for content delivery	Configures a Differentiated Services Code Point (DSCP) value for the QoS on a per-delivery service basis. If you choose **Other**, enter a decimal value in the corresponding field. **Note** This is a Release 2.3 feature and applies only to Windows Media Stream Streaming and Web engines.
Comments	Information about the delivery service.

Step 4 Click **Submit** to save the settings.

To delete a delivery service, from the Delivery Service Table page, click the **Edit** icon next to the delivery service you want to delete, and click the **Trash** icon in the task bar.

SE and Content Acquirer Assignment or Device Group and Content Acquirer Assignment

Step 5 through Step 8 use the Assign Service Engines option to describe the procedure of assigning the Service Engines to the delivery service and selecting one of them as the Content Acquirer. If you have device groups defined, you can use the Assign Device Groups option instead. To assign device groups, follow Step 5 through Step 8 and substitute Device Groups for each instance of Service Engines or SE.

Note Use either Assign Service Engines, or Assign Device Groups to assign Service Engines and select a Content Acquirer.

Step 5 From the left-panel menu, click **Assign Service Engines**. The Service Engine Assignment page is displayed (Figure 5-4).

Figure 5-4 Service Engine Assignment Page

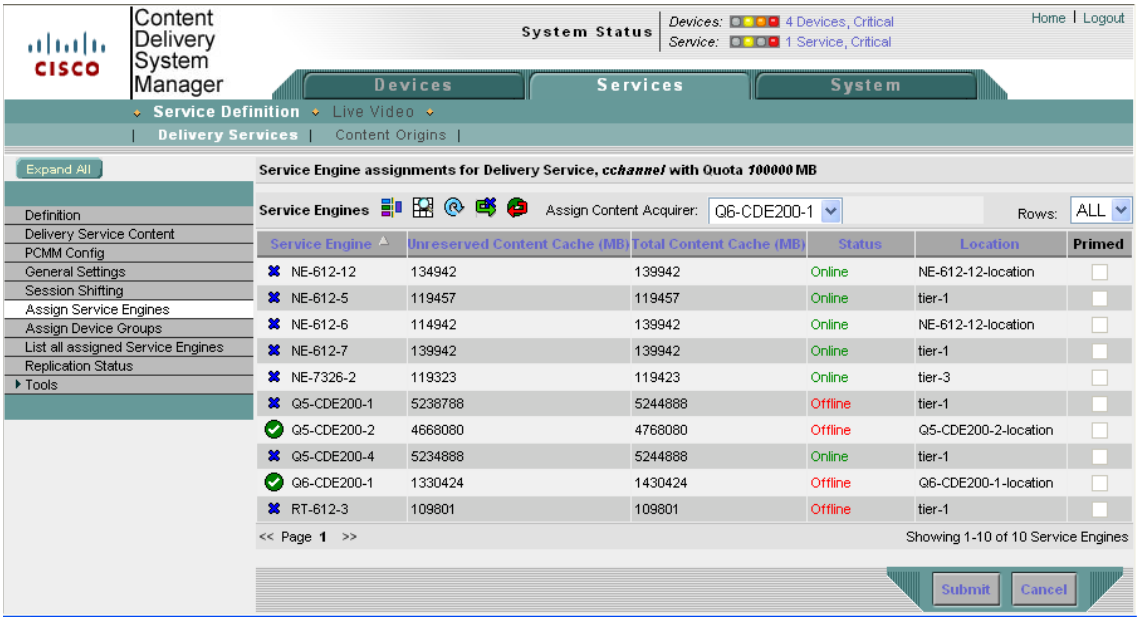

Step 6 Click the **Assign** icon (blue cross mark) next to the SE you want to assign to this delivery service. Alternatively, in the task bar, click **Assign All Service Engines**. The SE assignment states are described in Figure 5-5.

Figure 5-5 SE Assignment State

New Assign	Assigned and waiting for Submit	Assignment Submitted	Unassign Submitted Assign	Not modifiable. The quota on all the delivery services for this SE exceeds the disk space.
✖				✖

A green arrow wrapped around the blue cross mark indicates an SE assignment is ready to be submitted. To unassign an SE, click this icon.

Step 7 From the **Assign Content Acquirer** drop-down list in the task bar, choose an SE to be the Content Acquirer for this delivery service.

The list contains all SEs currently assigned to the delivery service.

The **Primed** check box indicates if an SE is primed with a live stream. For more information about priming, see the "Priming a Live Delivery Service" section on page 5-34.

Step 8 Click **Submit** to save the SE and Content Acquirer assignments.

A green circle with a check mark indicates an SE is assigned to this delivery service. To unassign the SE, click this icon, or click **Unassign All Service Engines** in the task bar. Click **Submit** to save the changes.

Note To view all the Service Engines assigned to the delivery service, in the left-panel menu, click **List assigned Service Engines**.

PCMM Configuration

Step 9 From the left-panel menu, click **PCMM Config**. The PCMM Configuration page is displayed (Figure 5-6).

Figure 5-6 PCMM Configuration Page

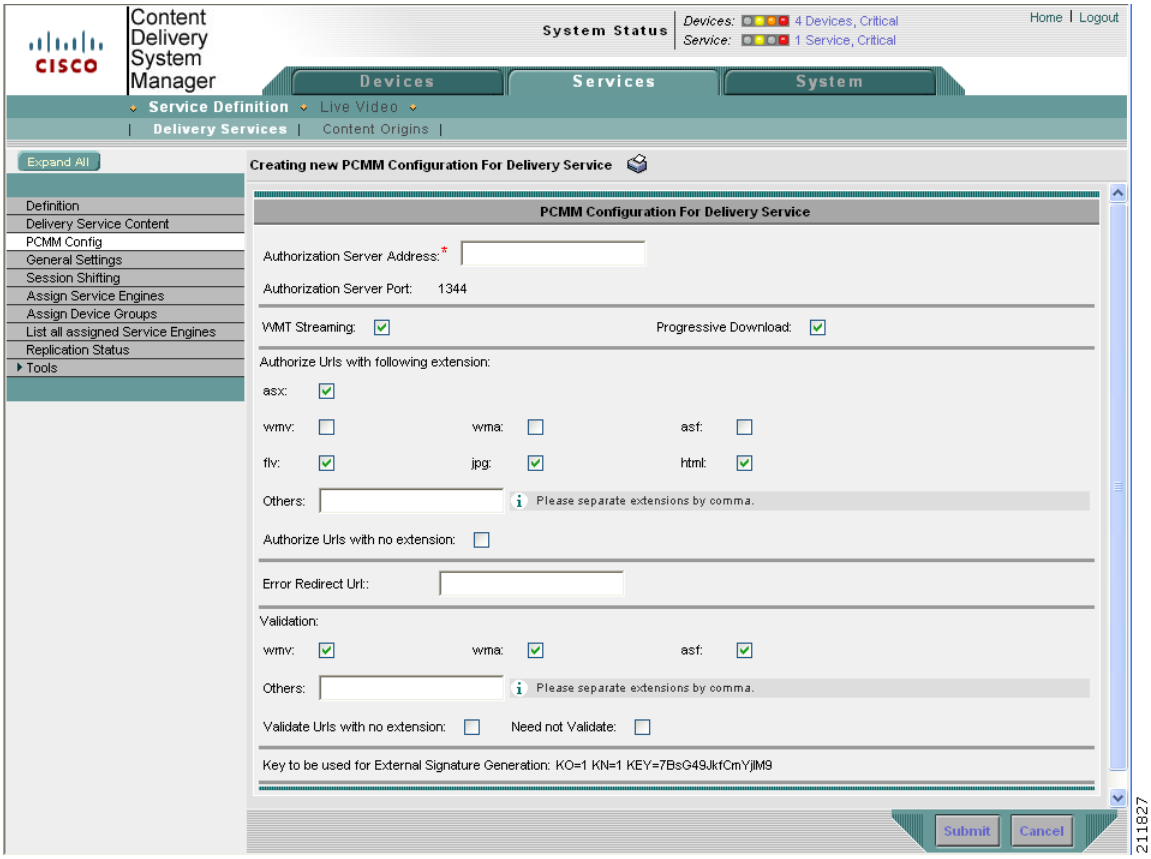

Step 10 Enter the settings as appropriate. See Table 5-3 for a description of the fields.

Table 5-3 PCMM Configuration Fields

Field	Description
Authorization Server Address	The IP address of the policy server.
Authorization Server Port (read only)	The port the policy server listens on for incoming authorization requests.

Table 5-3 PCMM Configuration Fields (continued)

Field	Description
WMT Streaming	When checked, an HTTP callout to the policy server is done for each stream request.
Progressive Download	When checked, an HTTP callout to the policy server is done for each download request.
Authorize URLs with following extensions	All URLs with file extensions that match those checked, or specified in the **Others** field (comma-separated), are authorized using ICAP.
Authorize URLs with no extensions	When checked, all URLs with no file extensions are authorized using ICAP.
Error Redirect URL	If the URL authorization fails, the request is redirected to the error URL specified in the field.
Validation	All URLs with file extensions that match those checked, or specified in the **Others** field (comma-separated), are validated in the SE.
Validation URLs with no extension	When checked, all URLs with no file extensions are validated in the SE.
Need not validate	When checked, validation is turned off.
Key to be used for External Signature Generation (Informational only)	The CDSM displays: `KO=1 KN=1 KEY=yEOCA6C7TWh4NAw9.` where • KO is the signature key owner. • KN is the signature key number. For information about configuring URL key signatures, see the "Configuring URL Signing" section on page 4-23.

Step 11 Click **Submit** to save the settings.

General Settings

Step 12 From the left-panel menu, click **General Settings**. The General Settings page is displayed.

Step 13 In the **Maximum Bitrate Limit per Session for HTTP** field, enter the maximum rate, in Kbps, at which a client can receive content.

This bit rate applies to content that is stored locally, specifically, prefetched, hybrid, or cached. For a cache miss, content is delivered at the rate the origin server sends it.

Step 14 In the **Maximum Bitrate Limit per Session for HTTP Cache Miss** field, enter the maximum rate, in Kbps, at which a client can receive content that was a cache miss.

This bit rate applies to content that is not stored locally, specifically, content requiring dynamic caching from the origin server or another SE in the delivery service that has the content.

Note This is a Release 2.1 feature.

Step 15 In the **Disable HTTP Download** field, check the check box to not allow clients to download HTTP content through this delivery service.

This option disables all HTTP-based content serving from this delivery service. The Web Engine will return a 403 forbidden message.

Note Because the Web Engine receives all HTTP requests before either Windows Media Streaming or Flash Media Streaming, if you disable HTTP download for a Windows Media Streaming delivery service or a Flash Media Streaming delivery service, and a client uses an HTTP request to download the SWF file, the Web Engine returns a 403 forbidden message. This is a Release 2.1 feature.

Step 16 Check the **Enable streaming over HTTP** check box and specify the file types in the **HTTP Allowed Extensions** field to configure progressive download or streaming for certain media files. This setting applies only to the following file types: .asf, none, .nsc, .wma, .wmv, and nsclog.

This delivery service setting has priority over the Windows Media Streaming engine settings on the Service Engines. If Windows Media Streaming is enabled on the Service Engines, and the media types are specified in the HTTP Allowed Extensions field, the delivery service streams the media types specified. If Windows Media Streaming is not enabled, or the media types are not specified in the HTTP Allowed Extensions field, the delivery service uses HTTP download.

Note This is a Release 2.3 feature.

Step 17 Click **Submit** to save the settings.

Session Shifting

For more information on session shifting, see the "3-Screen Session Shifting" section on page 1-5.

Note 3-Screen Session Shifting is a Release 2.2 feature and supports RTSP streaming for the Windows Media and Movie Streamer engines.

Step 18 From the left-panel menu, click **Session Shifting**. The Session Shifting page is displayed.

Step 19 Enter the settings as appropriate. See Table 5-4 for a description of the fields.

Table 5-4 Session Shifting Fields

Field	Description
Enable	When enabled, session shifting interoperates with the TV CDS.
TV-Streamer IP	The IP address of the TV Streamer participating in session shifting.
TV-Streamer Port	The port of the TV Streamer participating in session shifting.
Service Router	The Service Router that is designated as the centralized session manager. Only one Service Router is used for all session shifting delivery services. The same Service Router must be designated as the centralized session manager for all delivery services that use session shifting.

Table 5-4 Session Shifting Fields (continued)

Field	Description
Content File	For each session shifting file, click the **Browse** button to locate each XML file. The Choose File dialog box is displayed. Navigate to the file and click **Open**. To remove a file, click the **Trash** icon next to the file type. To view the file, click **List**. For more information on the session shifting XML files, see Appendix D, "Creating and Manipulating Session Shifting Files."
Subscriber File	
Profile File	

Step 20 Click **Submit** to save the settings.

Identifying Content

Content items are identified within the delivery service configuration for prefetch and hybrid ingests. Live program content is identified through the Live Program page, and therefore does not have content items listed for it in the delivery service.

Note The recommended maximum number of prefetched content items is 200,000.

When you configure a delivery service for content acquisition, you must choose one of the following methods:

- Identifying Content Using the CDSM

  The CDSM provides a user-friendly interface that you can use to add content items and specify crawl tasks without having to create and update a Manifest file. The CDSM automatically validates all user input and generates an XML-formatted Manifest file in the background that is free of syntax errors.

  Only one Manifest file is generated per delivery service for all content items. You can save your CDSM-generated Manifest file to any accessible location.

- Identifying Content Using a Manifest File

  The externally hosted Manifest files contain the XML tags, subtags, and attributes that define the parameters for content ingest. You must be familiar with the structure of the XML-based Manifest file and be sure the XML tags are properly formatted and syntactically correct before you can create and use Manifest files effectively.

Identifying Content Using the CDSM

There are several options in identifying content to be acquired using the CDSM. You can do any of the following:

- Identify a single content item.
- Define a crawl task that will search the origin server at the specified location (URL) and to the specified link depth, and create a list of all content that meets those specifications.
- Define a crawl task with the specifications described in the bullet above, and in addition specify content acquisition rules that further narrow the search.

- Select individual items by performing a quick crawl, and select the items from the crawl result list to be included in the content list.

For more information about the crawler feature, see the "Crawling" section on page 2-5.

To identify content for acquisition using the CDSM, do the following:

Step 1 Choose **Services > Service Definition > Delivery Services > Delivery Service Content**. The Content Table page is displayed with "Use GUI to specify content acquisition" as the method (Figure 5-7).

Figure 5-7 Content Table Page

Step 2 Click the **Add Content** icon in the task bar. The Content Manager page is displayed (Figure 5-8).

Figure 5-8 Content Manager Page

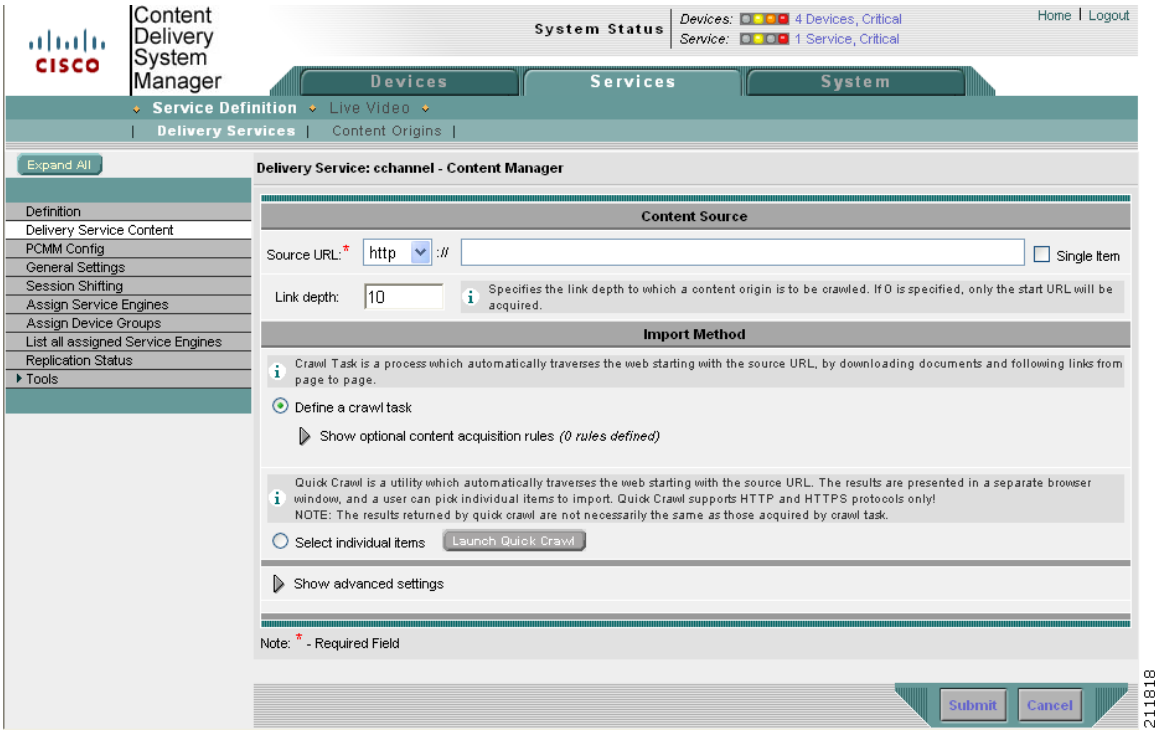
Content Delivery System Manager
System Status
Devices: 4 Devices, Critical
Service: 1 Service, Critical
Home | Logout
Devices
Services
System
Service Definition
Live Video
Delivery Services
Content Origins
Expand All
Definition
Delivery Service Content
PCMM Config
General Settings
Session Shifting
Assign Service Engines
Assign Device Groups
List all assigned Service Engines
Replication Status
Tools
Delivery Service: cchannel - Content Manager
Content Source
Source URL:*
http
://
Single Item
Link depth:
10
Specifies the link depth to which a content origin is to be crawled. If 0 is specified, only the start URL will be acquired.
Import Method
Crawl Task is a process which automatically traverses the web starting with the source URL, by downloading documents and following links from page to page.
Define a crawl task
Show optional content acquisition rules (0 rules defined)
Quick Crawl is a utility which automatically traverses the web starting with the source URL. The results are presented in a separate browser window, and a user can pick individual items to import. Quick Crawl supports HTTP and HTTPS protocols only!
NOTE: The results returned by quick crawl are not necessarily the same as those acquired by crawl task.
Select individual items
Launch Quick Crawl
Show advanced settings
Note: * - Required Field
Submit
Cancel
211818

To edit a content item, click the **Edit** icon next to the content. For more information about manipulating the content items in the Content Table page, see the "Configuring Proxy Server Settings" section on page 5-21.

Step 3 Choose a protocol from the **Source URL** drop-down list, and enter the source URL in the associated field.

The source URL is the origin server domain name or IP address, followed by a path, or path and filename, if applicable.

Note The URL format for Server Message Block (SMB) servers is: \\SMB server:port\sharedfolder\filepath. If port is not specified in the URL, the default port, 139, is used. Maximum file size, when using SMB for acquisition, is 2 GB. Symbolic links within exported file systems (SMB or NFS) must contain a relative path to the target file, or the target file should be copied into the exported volume.

Step 4 Do one of the following:

a. To identify a single content item, check the **Single Item** check box, and go to the "Configuring Advanced Settings" section on page 5-17 in this procedure.

b. To define a crawl, uncheck the **Single Item** check box, and in the **Link Depth** field, enter the depth of the links to search. Go to the "Defining a Crawl Task" section on page 5-14 in this procedure.

c. To perform a quick crawl, uncheck the **Single Item** check box, and, in the **Link Depth** field, enter the depth of the links to search. Go to the "Launching Quick Crawl" section on page 5-16 in this procedure.

The crawler feature starts with the Source URL, identifies every web link in the page, and adds every link to the list of URLs to search, until the links have been followed to the specified depth.

The Link Depth field specifies how many levels of a website to crawl or how many directory levels of an FTP server to search. This is optional. The range is –1 to 2147483636.

If the depth is –1, there is no depth constraint.

If the depth is 0, content is acquired only at the starting URL.

If the depth is 1, content is acquired starting at the URL and includes content the URL references.

Defining a Crawl Task

To define a crawl task, do the following:

Step 1 Click the **Define a Crawl Task** radio button.

Step 2 Do *one* of the following:

a. Click **Submit** (or **Update** if you are editing an existing content) to add a crawl task to the delivery service. The local Manifest file is automatically reparsed, changes are detected, and the corresponding content items are acquired or removed.

b. Go to the "Configuring Advanced Settings" section on page 5-17, if applicable.

c. Continue to the next step and create acquisition rules.

Step 3 Click the **Show Optional Content Acquisition Rules** arrow to further refine the crawl task. The fields in the acquisition rules are displayed (Figure 5-9), and the arrow becomes the **Hide Optional Content Acquisition Rules** arrow.

Figure 5-9 Content Manager Page—Acquisition Rules Fields

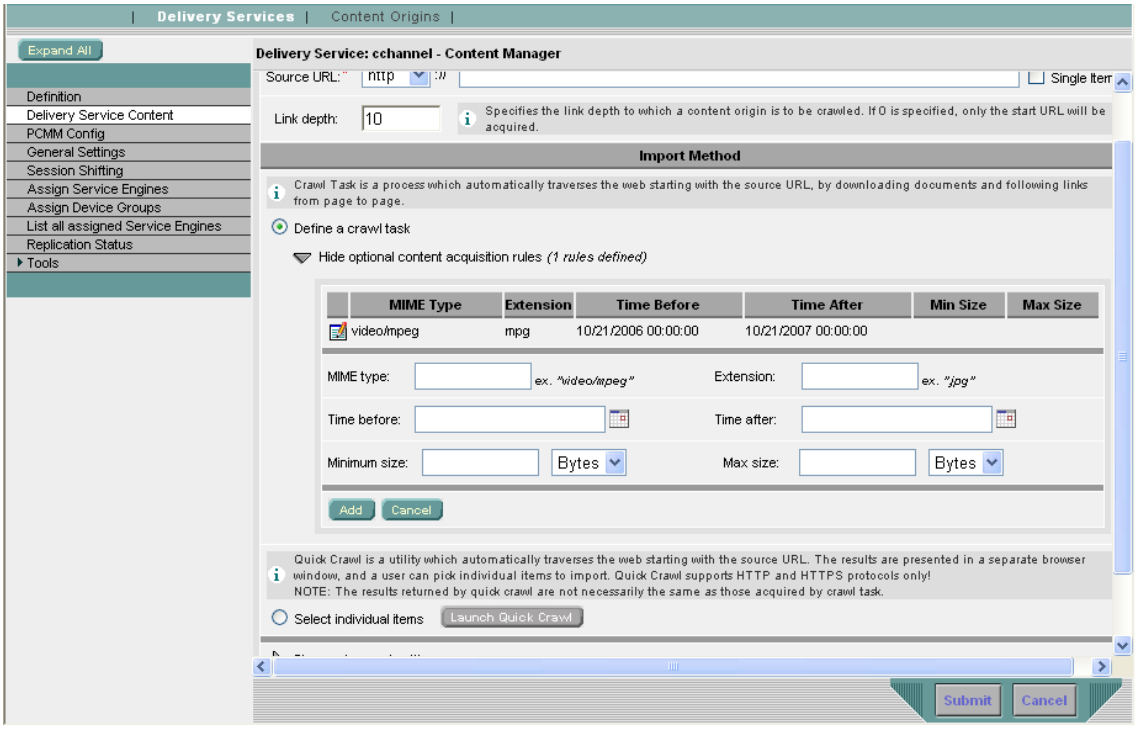
Delivery Services
Content Origins
Expand All
Delivery Service: cchannel - Content Manager
Definition
Delivery Service Content
PCMM Config
General Settings
Session Shifting
Assign Service Engines
Assign Device Groups
List all assigned Service Engines
Replication Status
Tools
Source URL:
http
Single Item
Link depth:
10
Specifies the link depth to which a content origin is to be crawled. If 0 is specified, only the start URL will be acquired.
Import Method
Crawl Task is a process which automatically traverses the web starting with the source URL, by downloading documents and following links from page to page.
Define a crawl task
Hide optional content acquisition rules (1 rules defined)
MIME Type
Extension
Time Before
Time After
Min Size
Max Size
video/mpeg
mpg
10/21/2006 00:00:00
10/21/2007 00:00:00
MIME type:
ex. "video/mpeg"
Extension:
ex. "jpg"
Time before:
Time after:
Minimum size:
Bytes
Max size:
Bytes
Add
Cancel
Quick Crawl is a utility which automatically traverses the web starting with the source URL. The results are presented in a separate browser window, and a user can pick individual items to import. Quick Crawl supports HTTP and HTTPS protocols only!
NOTE: The results returned by quick crawl are not necessarily the same as those acquired by crawl task.
Select individual items
Launch Quick Crawl
Submit
Cancel
211820

Step 4 Enter the settings as appropriate. See Table 5-5 for a description of the fields.

Table 5-5 Acquisition Rule Fields

Field	Description
MIME Type	A content item qualifies for acquisition only if its MIME type matches this MIME type (for example, video/mpeg).
Extension	A content item is acquired only if its extension matches this extension.
Time Before	Files that were modified before this time qualify for acquisition. Use the dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default. Alternatively, click the **Calendar** icon to choose a date from the calendar and enter a time, and click **Apply**.
Time After	Files that were modified after this time qualify for acquisition. Use the format dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default. Alternatively, click the **Calendar** icon to choose a date from the calendar and enter a time, and click **Apply**.
Minimum Size	Content equal to or larger than this value qualifies for acquisition. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.
Max Size	Content equal to or less than this value qualifies for acquisition. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.

Step 5 Click **Add** to add the rule to the rules list. An entry is added showing the values under each column heading.

Note A maximum of ten rules can be configured for each crawl task.

To modify a content acquisition rule, click the **Edit** icon next to the rule. Once you have finished, click the small **Update** button in the content acquisition rules section to save the edits.

To delete a content acquisition rule, click the **Edit** icon next to the rule. Click the **Delete** button in the content acquisition rules section. The rule is removed from the rules listing.

Step 6 When you have finished adding and modifying content acquisition rules, do one of the following:

a. If this is a new crawl task, click **Submit**.

b. If you are editing an existing crawl task, click **Update**.

c. Go to the "Configuring Advanced Settings" section on page 5-17, if applicable.

Launching Quick Crawl

Quick Crawl is a utility that automatically crawls websites starting from the specified source URL. You can use this utility when you know only the domain name and not the exact location of the content item. Quick Crawl supports crawling only for HTTP and HTTPS acquisition protocols.

To launch a quick crawl, do the following:

Step 1 Click the **Select Individual Items** radio button and click **Launch Quick Crawl**. The Quick Crawl Filter window is displayed.

Step 2 Enter the settings as appropriate. See Table 5-6 for a description of the fields.

Table 5-6 Quick Crawl Filter Fields

Field	Description
MIME Type	A content item is listed in the results only if its MIME type matches this MIME type (for example, video/mpeg).
Extension	A content item is listed only if its extension matches this extension.
Modified After	A content item is listed only if it was modified after this date. Click the **Calendar** icon to choose a date from the calendar, or enter the date in mm/dd/yyyy format.
Modified Before	A content item is listed only if it was modified before this date. Click the **Calendar** icon to choose a date from the calendar, or enter the date in mm/dd/yyyy format.
Minimum Size	Content equal to or larger than this value is listed in the results. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.
Max Size	Content equal to or less than this value is listed in the results. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.
Link Depth	How many levels of a website to crawl or how many directory levels of an FTP server to crawl. The range is –1 to 2147483636. If entered, the value from the Content Manager page is brought over to this field.

Table 5-6 Quick Crawl Filter Fields (continued)

Field	Description
Max Item Count	The maximum number of content items that is listed in the results. The maximum value is 1000.
Domain	The *host.domain* portion of the source URL. Edit this field to limit the search to a specific host on a domain.
Username	The username to log in to host servers that require authentication.
Password	The password for the user account.

Step 3 Click **Start Quick Crawl** to being search. The Searching for Content status displays a progress bar and shows the number of items found.

Click **Show Results** to display the content items before the search is complete.

Click **Refresh Results** to refresh the progress bar.

When finished, the search results list the MIME type, size, date modified, and URL of each content item that met the search criteria.

Step 4 Check the check box next to the content items you want to include in this delivery service. Use the **Row** drop-down list to show all content items, or use the **Page** option at the bottom of the table to go to the next page.

Alternatively, click **Select All** to select all content items. To deselect all, click **Select None**.

Step 5 Click **Add Selected** to add all selected content items to the delivery service. The Content Table page is displayed with all the selected content items listed.

Click **Show Filter** to return to the filter and change the filter settings.

Step 6 To configure advanced settings for the content items listed, click **All** at the bottom of the Content Table page, and then click **Edit Selected Items**. The Content Manager page is displayed with the Advanced Settings option.

Configuring Advanced Settings

Advanced settings offer controls on how the content is delivered to the client devices.

To configure the advanced settings, do the following:

Step 1 Click the **Show Advanced Settings** arrow. The Advanced Settings fields are displayed (Figure 5-10), and the arrow becomes the **Hide Advanced Settings** arrow.

Figure 5-10 Content Manager Page—Advanced Settings Fields

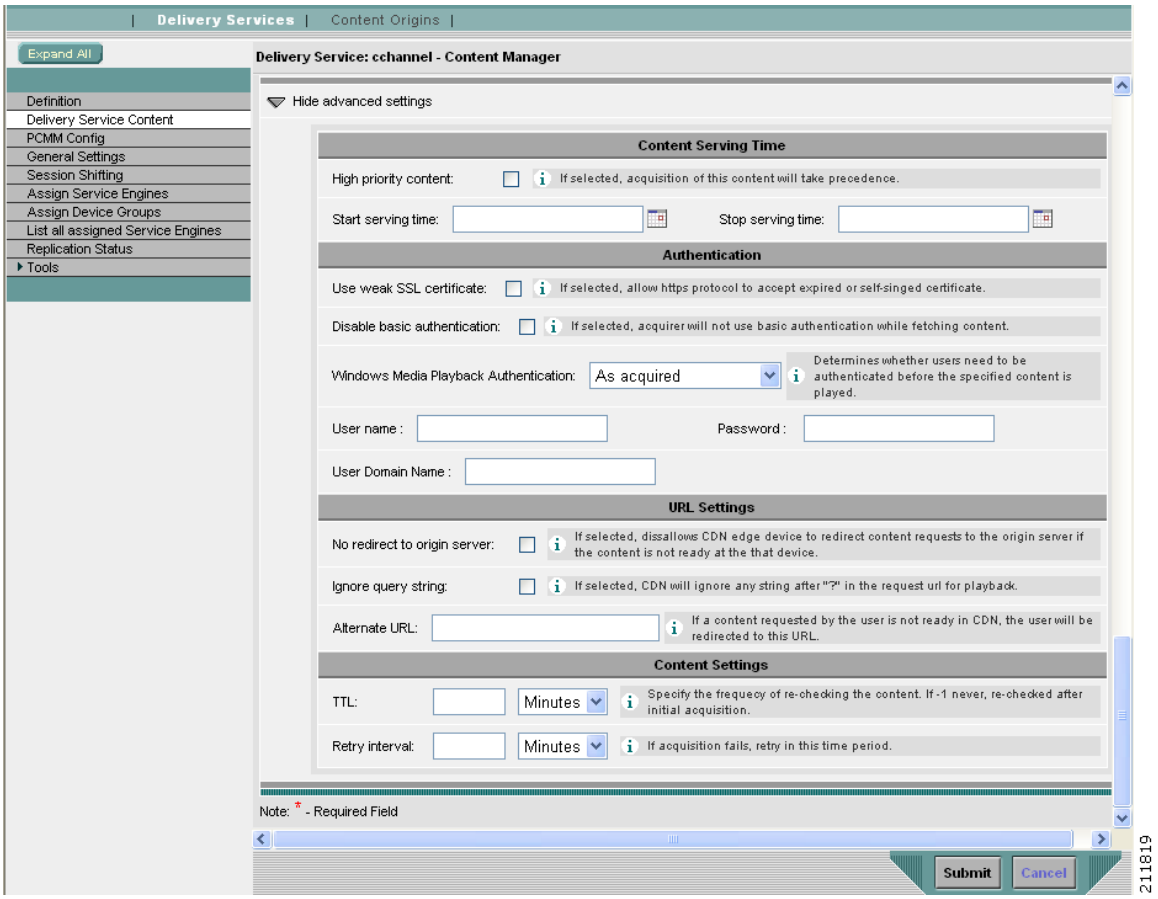
Delivery Services | Content Origins |
Expand All
Delivery Service: cchannel - Content Manager
Definition
Delivery Service Content
PCMM Config
General Settings
Session Shifting
Assign Service Engines
Assign Device Groups
List all assigned Service Engines
Replication Status
Tools
Hide advanced settings
Content Serving Time
High priority content:
If selected, acquisition of this content will take precedence.
Start serving time:
Stop serving time:
Authentication
Use weak SSL certificate:
If selected, allow https protocol to accept expired or self-singed certificate.
Disable basic authentication:
If selected, acquirer will not use basic authentication while fetching content.
Windows Media Playback Authentication:
As acquired
Determines whether users need to be authenticated before the specified content is played.
User name :
Password :
User Domain Name :
URL Settings
No redirect to origin server:
If selected, dissallows CDN edge device to redirect content requests to the origin server if the content is not ready at the that device.
Ignore query string:
If selected, CDN will ignore any string after "?" in the request url for playback.
Alternate URL:
If a content requested by the user is not ready in CDN, the user will be redirected to this URL.
Content Settings
TTL:
Minutes
Specify the frequecy of re-checking the content. If -1 never, re-checked after initial acquisition.
Retry interval:
Minutes
If acquisition fails, retry in this time period.
Note: * - Required Field
Submit
Cancel
211819

Step 2 Enter the settings as appropriate. See Table 5-7 for a description of the fields.

Table 5-7 Advanced Settings for Serving Content

Field	Description
Content Serving Time	
High Priority Content	Specifies the importance, and therefore the processing order, of the item acquisition or crawl task.
Start Serving Time	Specifies the time for the SE to start delivering content. Use the format dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default. Alternatively, click the **Calendar** icon to choose a date from the calendar and enter a time, and click **Apply**. If you do not specify a time, content is ready for delivery as soon as it is acquired and distributed to the SEs in the delivery service.

Table 5-7 Advanced Settings for Serving Content (continued)

Field	Description
Stop Serving Time	Specifies the time for the SE to stop delivering content. Use the dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default. Alternatively, click the **Calendar** icon to choose a date from the calendar and enter a time, and click **Apply**. If you do not specify a time, content continues to be available for delivery until you remove it from the delivery service either by changing the local Manifest file, using the Content Removal page, or renaming the delivery service. For information about the Content Removal page, see the "Removing Content" section on page 8-28.
Authentication	
Use weak SSL certificate	If checked, allows acceptance of expired or self-signed certificates during authentication.
Disable basic authentication	If checked, NTLM headers are not stripped off that would allow fallback to the basic authentication method while acquiring content.
Windows Media Playback Authentication	Sets the authentication for Windows Media playback to one of the following: • As acquired—Requires authentication on playback based on settings from origin server. • Require authentication—Requires authentication upon playback. • No authentication—Does not require authentication upon playback.
User Name	Name of the user for authentication.
Password	Password of the user for authentication.
User Domain Name	NTLM user domain name for the NTLM authentication scheme.
URL Settings	
No Redirect to Origin Server	**Note** For Releases 2.0 and 2.1 only. If checked, disallows redirecting content requests to the origin server. This attribute is a per-content object attribute, meaning that if the content has been removed, the redirection settings do not apply. However, if the Content Acquirer fails to acquire the content, then the settings apply.
Ignore Query String	If checked, ignores any string after the question mark (?) character in the requested URL for playback.
Alternate URL	**Note** This feature has been deprecated.
Content Settings	
TTL	Time period for revalidation of content. Select unit of measure from the drop-down list. If no TTL is entered, the content is fetched only once, and its freshness is never checked again.
Retry Interval	Time period in which the Content Acquirer can attempt to acquire the content again if the acquisition fails.

Step 3 Click **Submit** to process the content request. When you click **Submit**, the local Manifest file for this delivery service is automatically reparsed, changes are detected, and the corresponding items are acquired or removed. This action, however, does not trigger a recheck of all the content in the delivery service.

Content Table

The Content Table page (Figure 5-11) offers the following functions through the task bar:

Refresh Table	Add Content	Delete Selected Items	Manage Host and Proxy Settings	Save Settings Locally	Process Content Changes

The **Refresh Table** icon refreshes the content table.

The **Add Content** icon allows you to add content items by displaying the Content Manager page.

To delete a content item, check the check box next to each item you want to delete, and click the **Delete Selected Items** icon. To select all content items, click **All**. To deselect all content items, click **None**.

Figure 5-11 Content Table Page

For information on the **Manage Host and Proxy Settings** icon, see the "Configuring Proxy Server Settings" section on page 5-21.

After you save the CDSM-generated Manifest file by clicking **Submit** in the Content Manager page, you can save the Manifest file locally, and modify it. Choose the content item in the table, and click the **Save Settings Locally** icon in the task bar. A web browser window with the CDSM-generated Manifest file elements is displayed. Choose the **File Save As** option, enter a name for the Manifest file, and click **OK**. The Manifest file is saved on your PC. See Appendix B, "Creating Manifest Files," for more information.

To acquire configured content items immediately, click the **Process Content Changes** icon in the task bar.

Note If you change the Manifest file that you saved, and you want to use that Manifest file instead of the content that you defined in the CDSM, or if you want to use the Manifest file for another delivery service, then you must use the **Specify external manifest file** method and point to the Manifest file. When you

change the content acquisition method, any content items that you added are removed. For information about the Manifest file, see the "Identifying Content Using a Manifest File" section on page 5-22 and Appendix B, "Creating Manifest Files."

To edit multiple content items, check the check box next to each item you want to edit, and click **Edit Selected Items**.

Configuring Proxy Server Settings

When the Content Acquirer cannot directly access the origin server, because the origin server is set up to allow access only by a specified proxy server, you can configure acquisition through a proxy server. When a proxy server is configured for the Content Acquirer, the Content Acquirer contacts the proxy server instead of the origin server, and all requests to that origin server go through the proxy server.

Note Content acquisition through a proxy server is supported only for HTTP requests.

Note Before configuring a proxy server, verify that the Content Acquirer is able to ping the proxy server. If the proxy is not servicing the configured port, you will get the message: "failed: Connection refused."

To configure a proxy server for content items identified using the CDSM, do the following:

Step 1 From the Content Table page, click the **Manage Host and Proxy Settings** icon in the task bar.

The Content Hosts Table page is displayed, listing all previously created host URLs, the number of content items for each host, and a proxy server (if configured).

To return to the Content Table page, click **Return to Content Listing**.

Step 2 Check the check box next to each host you want to configure with a proxy server.

Step 3 Click **Manage Proxy for Selected Hosts**. The Proxy Server page is displayed.

Under the Defining Proxy Server for the Following Hosts heading, a bulleted list of host servers is displayed for which proxy servers are being configured.

Step 4 In the Proxy Server Specifications area, enter the settings as appropriate. See Table 5-8 for a description of the fields.

Table 5-8 Proxy Server Fields

Field	Description
Proxy Host	Hostname or IP address of the proxy server used by the Content Acquirer for content acquisition. When you use a domain name instead of an IP address, make sure that the domain name can be resolved by the DNS servers.
Proxy Port	Port number of the proxy server on which the Content Acquirer fetches content. The range is from 1 to 65535.

Table 5-8 Proxy Server Fields (continued)

Field	Description
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off that would allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method and the username and password information can be passed to the origin server in clear text with a basic authentication header.
User Name	Name of the user to be authenticated to fetch the content.
Password	Password of the user to pass authentication from the proxy.

Note If the specified proxy fails, the Content Acquirer, by default, contacts the origin server directly and tries to fetch the content.

Step 5 Click **Add** to add the proxy server.

To edit the proxy server settings, choose the proxy server from the Select a Proxy Server list, and click **Edit**. The values for the proxy server are displayed in the Proxy Server Specification section. Once you have finished modifying the settings, click **Update**.

To delete the proxy server settings, choose the proxy server from the Select a Proxy Server list, and click **Delete**.

Step 6 To assign the proxy server to the host or hosts listed on this page, choose a proxy server from the Select a Proxy Server list, and click **Save Assignment**. The Content Hosts Table page is displayed.

Identifying Content Using a Manifest File

The Manifest file provides information about the content to be prefetched, or fetched at a later time (as in hybrid ingest), or provides information about live content streamed through the delivery service.

Note Before configuring the CDSM to receive the Manifest file, you need to create one. See Appendix B, "Creating Manifest Files" for details on creating a Manifest file. After you create the Manifest file, use the Manifest Validator utility to verify the syntax. See the "Manifest Validator Utility" section on page B-14 for more information.

To configure the Manifest file settings, do the following:

Step 1 Choose **Services > Service Definition > Delivery Services > Delivery Service Content**. The Content Table page is displayed with Use GUI to specify content acquisition as the method.

Step 2 To change to the Specify external Manifest file method, do the following:

a. Click **Change Method**.

b. From the drop-down list choose, **Specify external manifest file**.

c. Click **Save**.

d. In the confirmation dialog box, click **OK**.

The Content Manager page displays the Manifest file settings (Figure 5-12).

Note When you change the content acquisition method from Use the GUI to specify content acquisition to Specify external Manifest file for an existing delivery service, any content items that you added using the CDSM are removed. To save the existing settings, click the **Save Settings Locally** icon in the task bar.

Figure 5-12 Content Manager Page—Manifest File Settings

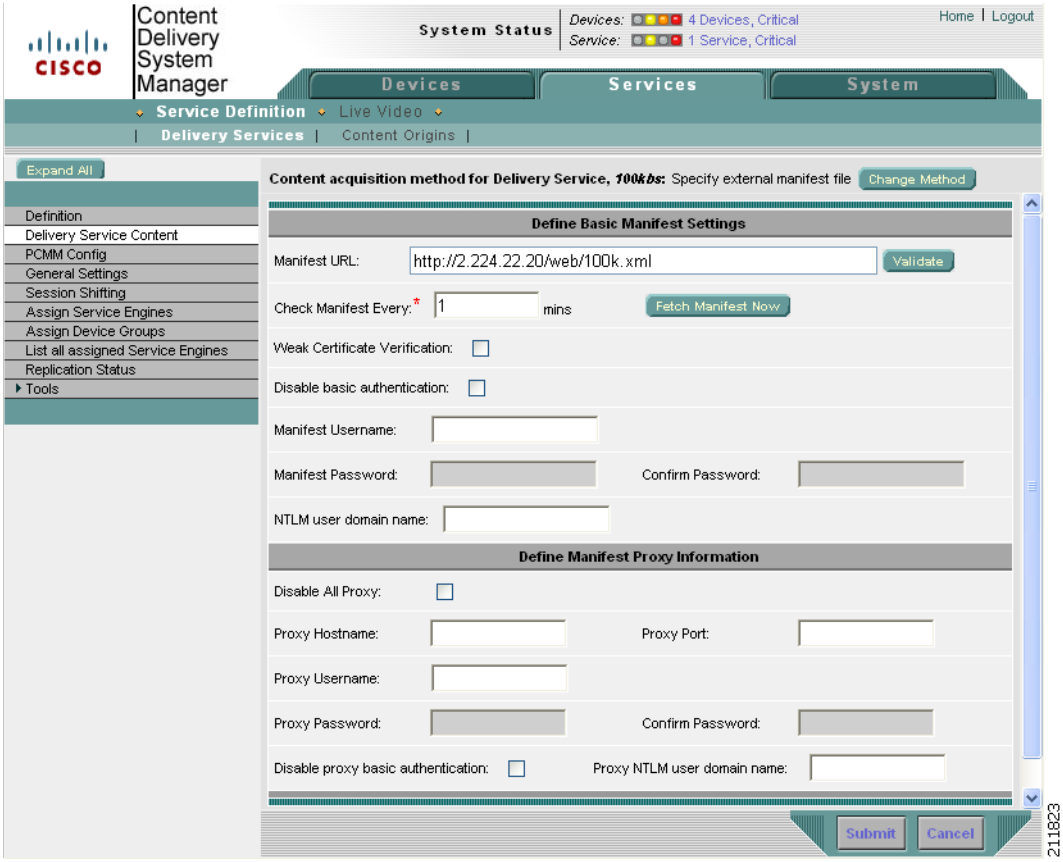

Step 3 Enter the settings as appropriate. See Table 5-9 for a description of the fields.

Table 5-9 Manifest File Settings Fields

Field	Description
Define Basic Manifest Settings	
Manifest URL	Address of the Manifest file for the delivery service. The Manifest URL must be a well-formed URL. If the protocol (FTP, HTTP, or HTTPS) for the URL is not specified, HTTP is used. To validate the Manifest file from this page, click **Validate**. A new window displays the validation results. For more information, see the "Manifest Validator Utility" section on page B-14.

Table 5-9 Manifest File Settings Fields (continued)

Field	Description
Check Manifest Every	Frequency in minutes (0 to 52560000) at which the Content Acquirer assigned to the delivery service checks for updates to the Manifest file. To fetch the Manifest file now, click **Fetch Manifest Now**.
Weak Certificate Verification	When checked, enables weak certificate verification for fetching the Manifest file. This is applicable when the Manifest file is fetched using HTTPS. **Note** To use weak certification for content ingest, you need to specify weak certification within the Manifest file.
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method, and the username and password information can be passed to the origin server in clear text with a basic authentication header.
Manifest Username	Username of the account that is allowed to fetch the Manifest file from the server. The Manifest username must be a valid ID. If the server allows anonymous login, the user ID can be null. **Note** The Manifest Username and Manifest Password fields allow you to enter any secure login information needed to access the Manifest file at its remote location.
Manifest Password	Password for the user.
Confirm Password	Password confirmation.
NTLM User Domain Name	NTLM user domain name to be allowed access by the NTLM authentication scheme configured on the server where the Manifest file is located.
Define Manifest Proxy Information	
Disable All Proxy	Disables the outgoing proxy server for fetching the Manifest file. Any outgoing proxy server configured on the Content Acquirer is bypassed, and the Content Acquirer contacts the server directly. See the "Configuring Web Engine HTTP Connections" section on page 4-43 for information about configuring outgoing HTTP proxy servers.
Proxy Hostname	Hostname or IP address of the proxy server used by the Content Acquirer to retrieve the Manifest file.
Proxy Port	Port number of the proxy sever where the Content Acquirer fetches the Manifest file. The range is from 1 to 65535.
Proxy Username	Name of the user to be authenticated to fetch the Manifest file.
Proxy Password	Password of the user to pass authentication on the proxy.
Confirm Password	Re-entry of the same password for confirmation to pass authentication on the proxy.

Table 5-9 Manifest File Settings Fields (continued)

Field	Description
Disable Proxy Basic Authentication	When checked, NTLM headers will not be stripped off to allow fallback to the basic authentication method against Microsoft Internet Information Services (IIS) servers.
Proxy NTLM User Domain Name	NTLM user domain name to be allowed access by the NTLM authentication scheme configured on the proxy.

Note When you configure a proxy server in the Manifest File Settings page, the proxy configuration is valid only for acquiring the Manifest file itself and not for acquiring the delivery service content. Requests for the Manifest file go through the proxy server, whereas requests for content go directly to the origin server.

Step 4 Click **Submit** to save the settings.

Proxy Server Settings

There are three ways to configure the proxy server when using a Manifest file to ingest content: through the CDSM, through the CLI, or through the Manifest file. If you need to configure the SE to use the proxy for both caching and prefetched content, use the CLI to configure the proxy. The CLI command is a global configuration command that configures the entire SE to use the proxy. If only the Content Acquirer portion of the SE needs to use the proxy for acquiring prefetched content, use the Manifest file to specify the outgoing proxy. When you configure the proxy server in the Manifest file, you are configuring the Content Acquirer to use the proxy to fetch content for the delivery service.

Note Proxy configurations in the Manifest file take precedence over proxy configurations in the CLI. Furthermore, a *noProxy* configuration in the Manifest file takes precedence over the other proxy server configurations in the Manifest file.

Configuring Programs

A program in the CDS is defined as a scheduled live or rebroadcast event that streams content to client devices. The CDS streams live or rebroadcast content by using the Movie Streamer or the Windows Media Engine. For more information, see the "Programs" section on page 2-9.

To view existing programs, see the "Viewing Programs" section on page 5-39.

Flash Media Streaming uses Real Time Media Protocol (RTMP) to stream live content by means of dynamic proxy. Configuration of live or rebroadcast programs is not required. When the first client requests live streaming content, the stream is created. For more information, see the "Live Streaming" section on page 1-18.

Caution If you have configured delivery services for live programs, make sure there are no external proxy servers physically located between your receiver SEs and your Content Acquirer that require proxy authentication. Also, make sure that proxy authentication is not enabled on any receiver SEs that might be in the logical, hierarchical path between the Content Acquirer and the receiver SE that is going to serve the live stream to the requesting clients. If a live stream encounters any device that requires proxy authentication, the stream will be dropped before it reaches its destination.

Note All SEs in a Windows Media live delivery service must have Real Time Streaming Protocol with TCP (RTSPT) enabled, because SEs must use the RTSPT protocol to communicate with each other. RTSPT is enabled by default.

Tip For information about testing a live or rebroadcast program, see Appendix G, "Testing the Internet Streamer CDS."

Note The following rules apply to live splitting for Movie Streamer:

1. For unicast streaming, the client request must be sent by means of RTSP.
2. For multicast streaming, the client request must be sent by means of HTTP.

Multicast Live Stream Interruptions

During a Windows Media live broadcast, any interruption of the live stream that lasts five minutes or longer causes the multicast broadcast to cease for the duration of the currently scheduled period. If the live stream is interrupted for less than five minutes, the broadcast resumes.

Live stream interruptions can be caused by unexpected encoder failures or by an operational restart. If the live stream stops for more than five minutes and resumes later while the program is still scheduled, you can modify the schedule or any other attribute of the program (such as the description) to trigger a restart of the multicast broadcast. Restarting might take up to five minutes under these circumstances.

This does not apply to unicast delivery of a Windows Media live event or to Movie Streamer live programs.

Defining a Program

To define a live or rebroadcast program, do the following:

Step 1 Choose **Services > Live Video > Live Programs**. The Program Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Program Definition page is displayed.

To edit an existing program, click the **Edit** icon next to the program name.

Step 3 In the **Name** field, enter a unique name for the program.

Step 4 From the **Type** drop-down list, choose a program type.

Step 5 Check the **Auto Deletion** check box if you want the program to be automatically deleted 24 hours after it has finished. This option only applies to live programs.

Note The Auto Deletion check box is not supported in Release 2.1 and subsequent releases.

Step 6 In the **Description** field, enter information about the program.

Step 7 Click **Submit** to save the settings.

You have defined the type of program that you want to configure. Proceed to the section for configuring that type of program:

- To configure Movie Streamer live and Windows Media live programs, see the next section, "Configuring Live Programs."
- To configure Windows Media rebroadcast and Movie Streamer rebroadcast programs, see the "Configuring a Rebroadcast" section on page 5-34.

For information about copying a program, see the "Copying a Program" section on page 5-41.

Configuring Live Programs

Once you have defined the program type, you must select a live delivery service, configure the streaming, and create a schedule. This procedure takes you through these steps and assumes you have already defined the program (see the "Defining a Program" section on page 5-27).

To configure a Movie Streamer live or Windows Media live program, do the following:

Step 1 After you have chosen a program from the Program Table page, click **Select Live Delivery Service**. The Select Live Delivery Service page is displayed listing all the live delivery services configured.

Step 2 Click the radio button next to the name of the live delivery service you want to associate with the program and click **Submit**. Alternatively, click the **Create New Live Delivery Service** icon in the task bar.

If you are creating a new live delivery service, the New Live Delivery Service page is displayed.

a. The **Name** field is automatically populated with a unique delivery service name. If you wish to change the name given by default, enter a unique name for the delivery service in this field.

b. From the **Content Origin** drop-down list, choose a Content Origin.

c. Click **Submit** to save the settings.

Step 3 From the left-panel menu, choose **Assign Service Engines**. The Service Engine Assignment page is displayed (Figure 5-13).

Figure 5-13 Service Engine Assignment Page

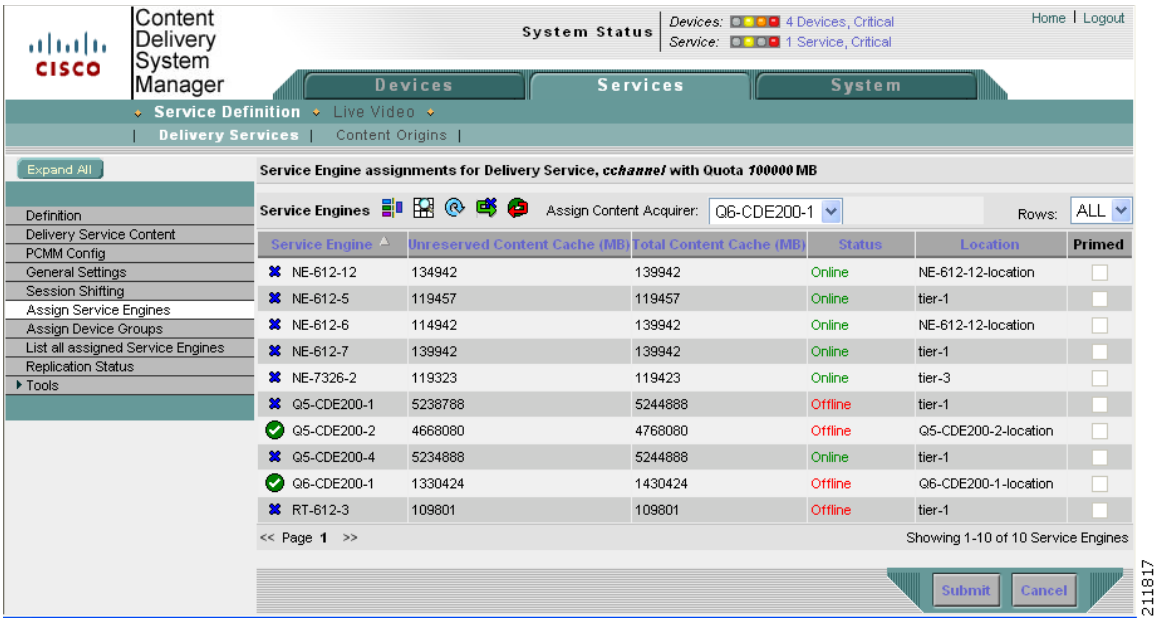

Step 4 Click the **Assign** icon (blue cross mark) next to the SE you want to assign to this delivery service. Or, in the task bar, click the **Assign All Service Engines** icon. The SE assignment states are described in Figure 5-14.

Figure 5-14 SE Assignment State

New Assign	Assigned and waiting for Submit	Assignment Submitted	Unassign Submitted Assignment	Not modifiable. The quota on all the delivery services for this SE exceeds the disk space

A green arrow wrapped around the blue cross mark indicates an SE assignment is ready to be submitted. To unassign an SE, click this icon.

Step 5 From the **Assign Content Acquirer** drop-down list in the task bar, choose an SE to be the Content Acquirer for this live delivery service.

The list contains all SEs currently assigned to the delivery service.

Step 6 Check the **Primed** check box for each SE you want to prime with the live stream. For more information about priming, see the "Priming a Live Delivery Service" section on page 5-34.

Step 7 Click **Submit** to save the SE and Content Acquirer assignments.

A green circle with a check mark indicates an SE is assigned to this delivery service. To unassign the SE, click this icon, or click the **Unassign All Service Engines** icon in the task bar. Click **Submit** to save the changes.

Step 8 In the left-panel menu, choose **Live Streaming**. The Live Stream Settings page is displayed.

The Live Stream Setting page differs depending on whether you are configuring a Movie Streamer live stream or a Windows Media live stream.

Step 9 Enter the settings as appropriate. See Table 5-10 for a description of the Windows Media Live Stream Settings fields, and Table 5-11 for a description of the Movie Streamer Live Stream Settings fields.

Table 5-10 Windows Media Live Stream Settings Fields

Field	Description
Live Source URL	The URL of the origin Windows Media encoder or Windows Media server using the following format: • http://*WMencoder_or_WMStreamerServer*:*port*/*path*/*file* • rtsp://*WMencoder_or_WMStreamerServer*:*port*/*path*/*file* For encoder failover, you can specify more than one encoder. Separate live source URLs in the list by using a semicolon (;). **Note** If you use a .wsx file as the Live Source URL and specify the encoders within the .wsx files, failover does not work for unicast-in multicast-out. We recommend you use a managed live-based encoder with redundancy, as it supports encoder failure with all type of streams.
Enable Unicast Delivery to Client	If enabled, the program uses unicast transmission.
Unicast URL Reference	If **Enable Unicast Delivery to Client** is checked, this field is auto-populated with a list of suggested URLs created from the Origin Server and the Service Routing Domain Name fields associated with the live delivery service. Choose one from the drop-down list.
Customized URL	**Note** This is a Release 2.0 feature. This field is available when the Customized Reference URL is selected in the Unicast URL Reference field. Use one of the following formats for the custom URL for unicast: rtsp://*OriginServerFQDN/customized-name* http://*OriginServerFQDN/customized-name*
Enable Multicast Delivery to Client	If enabled, the program uses multicast transmission.
Multicast URL Reference	If **Enable Multicast Delivery to Client** is checked, this field is auto-populated with a list of suggested URLs created from the Origin Server and the Service Routing Domain Name fields associated with the live delivery service. Choose one from the drop-down list.
NSC Reference for Multicast	The URL for the NSC file used for a server-side playlist as the media source in a multicast program. This field is available when **Enable Multicast Delivery to Client** is checked.

Table 5-10 Windows Media Live Stream Settings Fields (continued)

Field	Description
Customized URL	**Note** This is a Release 2.0 feature. This field is available when the Customized Reference URL is selected in the Multicast URL Reference field. Use the following format for the custom URL for multicast: **http://*any SEinLiveDeliveryService/customized-name.nsc***
Multicast Address and Port	The multicast address and port to use for streaming this program using multicast. The address range is 224.0.0.0 to 239.255.255.255. The port number must be even, and within the range of 1 to 65535. These values must be unique within the system. **Note** Auto Select is a Release 2.0 feature. Click **Auto Select** to select a multicast address from the multicast address pool. See the "Configuring a Multicast Address Pool" section on page 5-38 for more information.
Multicast TTL	Specify the multicast time to live (number of hops). The default is 15 hops.

Table 5-11 Movie Streamer Live Stream Settings Fields

Field	Description
Origin Server SDP File URL	The URL for the Session Description Protocol (SDP) file generated on the encoder. From the drop-down list, select either **rtsp** or **http**, and enter the remainder of the URL in the field. The remainder of the URL format is host [:port]/[filename], where the port and filename are optional. For the Darwin Streaming Server encoder, you need to specify the SDP file. For the Digital Rapid encoder, you do not need to specify the SDP file. In Release 2.2, when you click the **Auto Populate** button, the Incoming Live Streams Settings fields (the next section on the Live Streaming Settings page) are automatically populated based on the Origin Server SDP File URL **Note** In Releases 2.0 and 2.1, the Content Acquirer port must be zero (0) if the source is multicast push.
Backup SDP URL	The backup URL for the SDP file. This field is only for RTSP. Add a valid backup URL and click **Auto Populate**. The Incoming Live Streams Settings backup fields (the next section on the Live Streaming Settings page) are automatically populated based on the Backup SDP URL The Cisco CDS only supports failover between a primary Content Origin server and a backup Content Origin server for a Movie Streamer live program when the backup Content origin server uses the same codec as the primary. In Release 2.2, when you click the **Auto Populate** button, the Incoming Live Streams Settings fields (the next section on the Live Streaming Settings page) are automatically populated based on the Backup SDP File URL **Note** This is a Release 2.2 feature.
Incoming Live Streams Settings[1]	
Note Manually enter these fields when Auto Populate cannot parse the SDP URL.	
Primary Source Server	The stream source IP address.
Backup Source Server	The backup stream source IP address.
Primary Receiving IP	For RTSP, the Primary Receiving IP is the IP address of the Content Acquirer acting as the primary receiver. This is always unicast-in. For HTTP, the Primary Receiving IP is the multicast-in IP address used to broadcast the live stream.
Backup Receiving IP	For RTSP, the Backup Receiving IP is the IP address of the Content Acqurier acting as the backup receiver. Both the primary and backup Content Acquirer are located in the root location of the delivery service. For HTTP, the Backup Receiving IP is the multicast-in IP address used to broadcast the live stream.
Receiving Ports	Receiving ports are used to define each port related to audio and video streams.
Outgoing Live Streams Settings	

Table 5-11 Movie Streamer Live Stream Settings Fields (continued)

Field	Description
Unicast URL Reference	This field is auto-populated with a list of suggested URLs by using the Origin Server and the Service Routing Domain Name fields associated with the live delivery service. Choose one from the drop-down list.
Customized URL	**Note** This is a Release 2.0 feature. This field is available when the Customized Reference URL is selected in the Unicast URL Reference field. Use one of the following formats for the custom URL for unicast: rtsp://*ServiceRouterDomainName/program-name* http://*ServiceRouterDomainName/program-name* This URL points to a meta file (SDP) that is generated and resides on an external server.
Enable Multicast Delivery to Client	If enabled, the program uses multicast transmission. If you wish to enable support for Content Acquirer failover, you must check this check box. Content Acquirer failover for a live program works only when the incoming stream is a multicast stream.
Multicast URL Reference	This field is available if the **Enable Multicast Delivery to Client** check box is checked. Use the following format for the multicast URL reference (Announce URL): http://*sourceHost_or_FQDN/path/filename.sdp* This URL points to a meta file (SDP) that is generated and resides on an external server.
Multicast TTL	Specify the multicast time to live (number of hops). The default is 15 hops
Multicast Address	**Note** This is a Release 2.2 feature. The multicast address to use for streaming this program using multicast. The address range is 224.0.0.0 to 239.255.255.255. These values must be unique within the system.
Multicast Port	**Note** This is a Release 2.2 feature. The multicast port to use for streaming this program using multicast. The port number range is 1 to 65535. These values must be unique within the system.
Live Streams[2]	
Source Server	The stream source IP address. Click **Add New Live Stream** to add another live stream.

Table 5-11 Movie Streamer Live Stream Settings Fields (continued)

Field	Description
Content Acquirer Port	The port number on the Content Acquirer that will receive the stream. **Note** To activate failover support for the Content Acquirer, you must enter zero (0).
Multicast Address and Port	The multicast address and port to use for streaming this program using multicast. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system. **Note** Auto Select is a Release 2.0 feature. Click **Auto Select** to select a multicast address from the Multicast Address pool. See the "Configuring a Multicast Address Pool" section on page 5-38 for more information.

1. The Incoming Live Streams Settings section is new for Release 2.2.
2. The Live Streams section is for Releases 2.0 and 2.1. You can define up to ten live streams for each Movie Streamer live program. Movie Streamer programs can have more than one live stream (audio, video, and slides).

Step 10 Click **Submit** to save the settings.

Step 11 From the left-panel menu, choose **Schedule**. The Schedule page is displayed.

Step 12 Click the **Play Forever** radio button to have the program play continuously.

Alternatively, click the **Schedule Playtime** radio button to schedule up to ten different playtimes. The Playtime Editor is displayed in the page.

To edit an existing playtime, click the **Edit** icon next to the Initial Start Time.

To delete an existing playtime, click the **Trash** icon next to the Initial Start Time.

Step 13 Enter the settings for the playtime as appropriate. See Table 5-12 for a description of the fields.

Table 5-12 Playtime Fields

Field	Description
Start Playback on	The start date and time for the program.
UTC or SE (Local) Time	Which clock the start time should use, UTC or SE local.
Duration	The length of the program. In the drop-down list, choose minutes, hours, or days as the unit of time.
Repeat Frequency	The repeat frequency has the following options: • Do Not Repeat—Plays once. • Repeat Every—Repeats every so many days, hours, or minutes. • Repeat Weekly—Repeats at the same hour on the days you choose.
Repeat Forever Repeat Until	These fields display when **Repeat Every** or **Repeat Weekly** are chosen for Repeat Frequency. Repeat Forever repeats the program forever using the repeat frequency set in the previous fields. Repeat Until repeats the program based on the repeat frequency set in the previous fields and until the date and time specified in this field.

Step 14 Click **Submit** to save the settings.

Click **Add Playtime** to add additional playtimes to an existing schedule. The Playtime Editor is displayed in the page.

Priming a Live Delivery Service

The first client requesting a program often experiences the longest wait time for the program to begin playing. Users can experience long wait times because of the full RTSP negotiation that is required to pull the live stream from the source. Delays can also occur if the edge SE has not buffered enough stream data to fill the media player's buffer at the time the program is requested. For Windows Media streaming, when the buffer is not filled, some data to the client might be sent at the suboptimal line rate instead of at the Fast Start rate.

Delivery services for unicast-managed live programs can be primed for faster start-up times. When a live delivery service is primed, a unicast-out stream is pulled from the origin server to an SE before a client ever requests the stream. When the first request for the stream goes out, the stream is already in the delivery service.

Note It is not possible to monitor non-primed streams because they are played directly from the Content Origin server. Primed streams can be monitored because they are buffered on the SE.

Configuring a Rebroadcast

Once you have defined the program type, you need to select media files, configure the streaming, and create a schedule. This procedure takes you through these steps and assumes you have already defined the program (see the "Defining a Program" section on page 5-27).

Note For rebroadcast programs, media can only be selected from one delivery service. The SEs and device groups assigned to the delivery service are selected automatically when you choose the media files for the program.

To configure a Movie Streamer rebroadcast or Windows Media rebroadcast program, do the following:

Step 1 After you have chosen a program from the Program Table page, click **Select Media**. The Select Media page is displayed.

Step 2 Choose a delivery service from the list by clicking the radio button next to the name of the delivery service and click **Show Media in Selected Delivery Service**. The Media File Selection pane is displayed.

Step 3 In the **Criteria** field, enter the search criteria for the media files you want to add to the program and click **Use Criteria**. All the media files that match the search criteria are displayed.

Use an asterisk (*) to match any number of characters, or a question mark (?) to match exactly one character. For example, use "*.mpg" for all files with the suffix "mpg," and "file?.mpg" to match file1.mpg, file2.mpg, and so on.

To start a new search, click **Select Media**.

To choose a new delivery service to choose files from, click **All Delivery Services**, choose a delivery service, and click **Show Media in Selected Delivery Service**.

Step 4 Check the **Pick** check box next to each file you want to rebroadcast and click **Add Media**. The files are displayed in the Media Files in Program pane.

To select all files, click **All**. To deselect all files, click **None**. The file list can span several pages. To see the files from the other pages, click the page number, or from the **Row** drop-down list, select one of the options.

Step 5 In the Media Files in Program pane, use the Up arrow and Down arrow next to each file to alter the order of the files. Files are played in the order in which they are listed.

Note The Up arrow and Down arrow are only displayed if the list of media files in the program is sorted by position. If you sort the media files by name or length, the arrows are not displayed.

Note Only one media file can be selected for Movie Streamer rebroadcasts.

To remove a media file from the list, check the **Pick** check box next to the file, and click **Remove Media**. To select all files, click **All**. To deselect all files, click **None**.

Step 6 Click **Submit** to save the settings.

Note For rebroadcast programs, media can only be selected from one delivery service. The SEs assigned to that delivery service are selected automatically when you choose the media files for the program. If at a later time you add new SEs to the delivery service, you must manually add them to the program.

Step 7 To add new SEs to the rebroadcast program, from the left-panel menu click **Assign Service Engines**. The Service Engine Assignment page is displayed.

Step 8 Click the **Assign** icon (blue cross mark) next to the SE you want to assign to this delivery service. Or, in the task bar, click the **Assign All Service Engines** icon. The SE assignment states are described in Figure 5-15.

Figure 5-15 SE Assignment State

New Assign	Assigned and waiting for Submit	Assignment Submitted	Unassign Submitted Assignment	Not modifiable. The quota on all the delivery services for this SE exceeds the disk space

A green arrow wrapped around the blue X indicates an SE assignment is ready to be submitted. To unassign an SE, click this icon.

Step 9 From the **Assign Content Acquirer** drop-down list in the task bar, choose an SE to be the Content Acquirer for this rebroadcast delivery service.

The list contains all SEs currently assigned to the delivery service.

Step 10 Check the **Primed** check box for each SE you want to prime with the rebroadcast stream. For more information about priming, see the "Priming a Live Delivery Service" section on page 5-34.

Step 11 Click **Submit** to save the SE and Content Acquirer assignments.

A green circle with a check mark indicates an SE is assigned to this delivery service. To unassign the SE, click this icon, or click the **Unassign All Service Engines** in the task bar. Click **Submit** to save the changes.

Step 12 From the left-panel menu, choose **Streaming**. The Streaming Settings page is displayed.

Step 13 Enter the settings as appropriate. See Table 5-13 for a description of the Windows Media Rebroadcast Stream Settings fields, and Table 5-14 for a description of the Movie Streamer Rebroadcast Stream Settings fields.

Table 5-13 Windows Media Rebroadcast Stream Setting s Fields

Field	Description
Multicast URL Reference	This field is auto-populated with a list of suggested URLs by using the Origin Server and the Service Routing Domain Name fields associated with the rebroadcast. Choose one from the drop-down list.
NSC Reference for Multicast	The URL for the NSC file used for a server-side playlist as the media source in a multicast program. This field is available when the Customized Reference URL is selected in the Multicast URL Reference field.
Customized URL	**Note** This is a Release 2.0 feature. This field is available when the Customized Reference URL is selected in the Multicast URL Reference field. Use the following format for the custom URL for multicast: http://*anySEinDeliveryService/program-name.nsc*
Multicast Address and Port	The multicast address and port to use for streaming this program using multicast. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system. **Note** Auto Select is a Release 2.0 feature. Click **Auto Select** to select a multicast address from the multicast address pool. See the "Configuring a Multicast Address Pool" section on page 5-38 for more information.
Multicast TTL	Specify the multicast time to live (number of hops). The default is 15 hops.

Table 5-14 Movie Streamer Rebroadcast Stream Settings Fields

Field	Description
Multicast URL Reference	This field is auto-populated with a list of suggested URLs by using the Origin Server and the Service Routing Domain Name fields associated with the rebroadcast. Choose one from the drop-down list. **Note** The Content Acquirer port must be zero (0) if the source is multicast push.
Customized URL	**Note** This is a Release 2.0 feature. This field is available when the Customized Reference URL is selected in the Multicast URL Reference field. Use one of the following formats for the custom URL for unicast: http://*sourceHost_or_FQDN/path/filename.sdp* rtsp://*ServiceRouterDomainName/program-name* http://*ServiceRouterDomainName/program-name* This URL points to a meta-file (SDP) that is generated and resides on an external server.
Multicast TTL	Specify the multicast time to live (number of hops). The default is 15 hops
Multicast Address and Port	The multicast address and port to use for streaming this program using multicast. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system. **Note** Auto Select is a Release 2.0 feature. Click **Auto Select** to select a multicast address from the multicast address pool. See the "Configuring a Multicast Address Pool" section on page 5-38 for more information. **Note** Because Movie Streamer rebroadcast files can contain multiple tracks (1 to 3), you can define up to three multicast addresses and ports for each track in the file. Click **Add Multicast Address/Port** to add another multicast address pool.

Step 14 Click **Submit** to save the settings.

Step 15 From the left-panel menu, choose **Schedule**. The Schedule page is displayed.

Step 16 Click the **Loop Back Continuously** radio button to have the program play continuously.

Alternatively, click the **Schedule Playback** radio button to schedule up to ten different playback times. The Playtime Editor is displayed in the page.

To edit an existing playtime, click the **Edit** icon next to the Initial Start Time.

To delete an existing playtime, click the **Trash** icon next to the Initial Start Time.

Step 17 Enter the settings for the playtime as appropriate. See Table 5-15 for a description of the fields.

Table 5-15 Playtime Fields

Field	Description
Start Playback on	The start date and time for the program.
UTC or SE (Local) Time	Which clock the start time should use, UTC or SE local.
Duration	The length of the program. In the drop-down list, choose minutes, hours, or days as the unit of time.
Playback Options	The playback options are the following: • Playback Once and Stop • Loop for number of minutes, hours, or days
Repeat Frequency	The repeat frequency has the following options: • Do Not Repeat—Plays once. • Repeat Every—Repeats every so many days, hours, or minutes. • Repeat Weekly—Repeats at the same hour on the days you choose.
Repeat Forever Repeat Until	These fields display when **Repeat Every** or **Repeat Weekly** are chosen for Repeat Frequency. Repeat Forever repeats the program forever using the repeat frequency set in the previous fields. Repeat Until repeats the program based on the repeat frequency set in the previous fields and until the date and time specified in this field.

Step 18 Click **Submit** to save the settings.

Click **Add Playtime** to add additional playtimes to an existing schedule. The Playtime Editor is displayed in the page.

Configuring a Multicast Address Pool

The multicast delivery feature is enabled by setting up a multicast address for a live or rebroadcast program to which different client devices, configured to receive content from the same program, can subscribe. The delivering device sends content to the multicast address set up at the SE, from which it becomes available to all subscribed receiving devices.

A set of multicast addresses can be specified either in the Program API or by using the CDSM. When a program requires a multicast address, you can specify the multicast address within the stream settings of the program, or you can have the CDSM select one of the addresses from the multicast address pool. Addresses are allocated for the life of a program.

When you request a specific address or a set of addresses to be used for a program, the CDSM issues only those addresses that are not used by any existing programs. If no addresses are available from the pool, or if the multicast pool has not been configured, users receive an error message.

To configure a pool of multicast addresses to be used for programs, do the following:

Step 1 Choose **Services > Live Video > Multicast Addresses**. The Multicast Addresses page is displayed.

Step 2 In the **Start Address** field, specify the first multicast IP address in the pool.

The range is 224.0.0.0 to 239.255.255.255.

Step 3 In the **End Address** field, specify the last multicast IP address in the pool.

The range is 224.0.0.0 to 239.255.255.255.

Step 4 In the **TTL field**, specify the time-to-live (number of hops) for all addresses configured in the pool.

The range is 1 to 255.

Step 5 Click **Submit** to save your settings.

The list of multicast addresses that have been currently configured for specific programs is displayed in the Multicast Address table. The User Specified column displays *true* if the user has already specified the particular address for a program.

Viewing Programs

The Programs Table page lists all of the programs defined in your CDS network. Programs can be defined through the CDSM or through an API.

The Programs Table page allows you to view scheduled programs by day, week, month, or year. You can sort and filter programs by name, type, or schedule. You can also preview live programs while they are playing. See the "Previewing a Program" section on page 5-41 for more information.

To view all the programs defined in your CDS network, follow these steps:

Step 1 Choose **Services > Live Video > Live Programs**. The Programs Table page displays with a list of all the programs that have been defined through either the CDSM or the Program API.

Step 2 Click the **Day**, **Week**, **Month**, or **Year** tab to view the playback schedules. Scheduled programs are listed by start time (initial start time plus any repeat intervals). Times begin with the current device time (current system time plus device time zone offset).

The **Unscheduled** tab displays all unscheduled programs defined in your CDS network. The **All** tab displays all the programs defined in your CDS network. The Programs Table page opens to the All view by default.

Step 3 Sort columns by clicking the column heading. You can also combine filtering conditions. For example, you can filter only Windows Media live programs and then choose the **Week** tab to view the week of November 23 to November 29, 2007. Table 5-16 describes the information that is displayed in this page.

Table 5-16 Programs Table Page Information

Item	Description
Tabs	
Day/Week/Month/Year	Lists programs based on their schedule. The current day, week, month, or year is displayed by default. You can navigate to the next or previous day, week, month, or year by clicking the back or forward arrows on either side of the date.
Unscheduled	Lists only programs with no schedule defined.
All	Lists all programs. This is the default view.

Table 5-16 Programs Table Page Information (continued)

Item	Description
Program Listing Table	
Program	Program name, which must be unique to the CDSM.
Type	Program type. Program types are: • Movie Streamer live • Movie Streamer rebroadcast • Windows Media live • Windows Media rebroadcast
Schedule	Describes the schedule. Options are: • None (the program has no schedule) • Loop continuously • Number of playtimes (the number of times that the program is scheduled to be shown) Start Time—Program start time in a scheduled view (Day, Week, Month, or Year tab). Lists up to three start times if repeat broadcasts are configured. Duration—Duration of the program or the looping time in a scheduled view (Day, Week, Month, or Year tab).

Viewing and Modifying API Programs

Programs created through APIs are based on a program file. A *program file* contains the elements that define the schedule, content, and presentation parameters. It is a text file written in XML format, similar to the Manifest file. The program file contains most of the program settings and resides on an external server. The CDSM gets the program file, parses it, and saves the program file to the database. The program is automatically updated at intervals by the CDSM refetching the program file and reparsing it. The program file supports RTSP.

In contrast, programs defined using the CDSM are not based on a program file; instead, the settings entered in the CDSM are saved directly to the database.

Programs created using an API can be viewed in the CDSM as read-only, and modifications to API programs can be done through the API. You can also edit the API program using the CDSM; however, if you choose this option, the information about the API program file is deleted and the program can no longer be modified through the API. A third option is to copy the API program using the CDSM Copy Program feature. The new copy will not contain the program file information and will be treated as a CDSM-generated program for the purposes of editing. (See the "Copying a Program" section on page 5-41.)

You can delete any program from the list (whether created through the CDSM or through an API) in the Programs Table page.

Previewing a Program

You can preview live programs by live split or by joining a multicast broadcast. Live programs can only be viewed during the scheduled playtime. You can preview a rebroadcast program by joining the multicast broadcast during the scheduled playtime.

To preview a live Movie Streamer or Windows Media program or scheduled rebroadcast, follow these steps:

Step 1 Choose **Services > Live Video > Live Programs**. The Programs Table page is displayed.

Step 2 Click the **Day**, **Week**, **Month**, or **Year** tab.

Step 3 Click the **Play** icon next to the name of a program. A program preview window pops up, displaying the program information with links to view the program.

Note The **Play** icon only appears while the live program is playing. If a program is not currently playing, you cannot view it.

Step 4 Click the URL reference link for the program. You have the option to choose a multicast or unicast URL reference, if such are defined for the program. A new window with the URL reference opens.

To successfully view the program, you must meet these conditions:

- You must be able to access the client network.
- You must have a Windows Media plug-in installed to view Windows Media live programs.
- You must have a QuickTime plug-in installed to view Movie Streamer live programs.

Copying a Program

The copy program feature allows you to create a copy of an existing program and then modify a subset of attributes, which eliminates the need to re-enter all the program settings each time you create programs with similar characteristics.

When you copy a program, a duplicate of the program is created and saved to the database. Any changes that you make to the new copy of the program do not affect the original program and vice versa. Note, however, that if multicast is configured, the multicast address and port cannot be copied. These parameters must be unique across the system. If a program address pool is configured, these parameters can be automatically selected by the system.

To create a copy of an existing program, follow these steps:

Step 1 Choose **Services > Live Video > Live Programs**. The Programs Table page is displayed.

Step 2 Click the **Edit** icon next to the name of the program that you want to copy. The Program Definition page is displayed.

Step 3 Click the **Copy Program** icon in the task bar. You are prompted to confirm your decision. Click **OK**. The window refreshes, displaying ProgramName_dup in the Name field.

Step 4 Edit any program information that you want to change. (See the "Defining a Program" section on page 5-27.)

Note You cannot change the program type.

Step 5 Click **Submit** to save the settings.

Step 6 Edit any of the other program properties found in the left-panel menu, such as the program schedule, program, or device assignments.

CHAPTER 6

Configuring the System

This chapter provides information on configuring the system parameters of the CDS. This chapter has the following major topics:

- Configuring AAA, page 6-1
- Changing a Password, page 6-5
- Configuring System Settings, page 6-6

Configuring AAA

Authentication determines who the user is and whether that user should be allowed access to the network or a particular device. It allows network administrators to bar intruders from their networks. It may use a simple database of users and passwords. It can also use one-time passwords.

Authorization determines what the user is allowed to do. It allows network managers to limit which network services are available to different users.

Accounting tracks what users did and when they did it. It can be used for an audit trail or for billing for connection time or resources used (bytes transferred).

Collectively, authentication, authorization, and accounting are sometimes referred to as AAA. Central management of AAA means the information is in a single, centralized, secure database, which is much easier to administer than information distributed across numerous devices.

In the CDS network, login authentication and authorization are used to control user access and configuration rights to the CDSM, SEs, and SRs. There are two levels of login authentication and authorization:

- Device
- CDSM

In a CDS network, user accounts can be created for access to the CDSM and, independently, for access to the SEs and SRs that are registered to the CDSM.

This section covers login authentication and authorization for the CDSM. For information about device login authentication and authorization, see the "Login Access Control" section on page 4-49 and the "Authentication" section on page 4-53.

Login authentication is the process by which CDSM verifies whether the person who is attempting to log in has a valid username and password. The person logging in must have a user account registered with the device. User account information serves to authorize the user for login and configuration

privileges. The user account information is stored in the AAA database. When the user attempts to log in, the CDSM compares the person's username, password, and privilege level to the user account information that is stored in the database.

Each user account can be assigned to a role and a domain. A *role* defines which CDSM configuration pages the user can access and which services the user has authority to configure or modify. A *domain* defines which entities in the network the user can access and configure or modify. You can assign a user account to zero or more roles, and to zero or more domains.

Creating, Editing, and Deleting Users

Note This section is addressed to users with administrator-level privileges (admin users) only.

Two default user accounts are preconfigured in the CDSM. The first account, called *admin*, is assigned the administrator role that allows access to all services and access to all entities in the system. This account cannot be deleted from the system, but it can be modified. Only the username and the role for this account are unchangeable. To change the password for this account, use the **username** *admin* **password** <*password*> command through the CLI.

The second preconfigured user account is called *default*. Any user account that is authenticated but has not been registered in the CDSM gets the access rights (role and domains) assigned to the default account. This account is configurable, but it cannot be deleted nor can its username be changed.

When you create a new user account in the CDSM, you have the option to create the user account in the CLI for the CDSM device at the same time. Using this option to create the new account in the CLI provides the following benefits:

- The user account is created in the primary and standby CDSM management databases and in the CDSM CLI from one central point.
- Users can change their passwords, and the password changes will propagate to a standby CDSM.

If you choose to create the user account from the CDSM *without* creating the user account in the CDSM CLI at the same time, the following results apply:

- The user account is created in the primary and standby CDSM management databases.
- No user account is created in the CDSM CLI, and the user *cannot* log in to the CDSM until an account is created from the CLI.
- Local users cannot change their passwords using the CDSM.
- Local users can change their passwords using the CLI; however, the password changes are not propagated from the CLI to the CDSM databases when the CLI user option is enabled in the CDSM.

  If a user account has been created from the CLI only, when you log in to the CDSM for the first time, the Centralized Management System (CMS) database automatically creates a user account (with the same username as configured in the CLI) with default authorization and access control. However, to change the password in this scenario, the user account must be explicitly configured from the CDSM with the CLI user option enabled.

To create or edit a user account, do the following:

Step 1 Choose **System > AAA > Users**. The User Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The User Account page is displayed.

To edit an account, click the **Edit** icon next to the username.

Note The User Account page can only be accessed by users with administrator-level privileges.

Step 3 In the **Username** field, enter the user account name.

Step 4 If you want to create a local user account with a password and privilege level from the CDSM, check the **Create CLI User** check box. The user account is created automatically in the CLI. To prevent the creation of a CLI user account from the GUI, leave the check box unchecked.

Step 5 In the **Password** field, enter a password for the CLI user account, and re-enter the same password in the **Confirm Password** field.

The password strength must be a combination of alphabetic character, at least one number, at least one special character, and at least one uppercase character.

Step 6 From the Privilege Level drop-down list, choose a privilege level for the CLI user account. The choices are 0 (zero) (normal user) or 15 (superuser). The default value is 0.

Note A superuser can use privileged-level EXEC commands, whereas a normal user can use only user-level EXEC commands.

Step 7 In the Username Information area, enter the following information about the user: First Name, Last Name, Phone Number, Email Address, Job Title, and Department.

Step 8 In the **Comments** field, enter any additional information about this account.

Step 9 Click **Submit** to save the settings.

Step 10 From the left-panel menu, click **Role Management**. The Role Management Table page is displayed.

To add roles, see the "Creating, Editing, and Deleting Roles" section on page 6-4.

To view the setting for the role, click the **View** (eyeglasses) icon next to the role.

Step 11 Click the **Assign** icon next to each role name you want to assign to the user account.

To remove the role from the user account, click the **Assign** icon again.

To assign all roles, click the **Assign All** icon in the task bar. To unassign all roles, click the **Remove All** icon in the task bar.

Step 12 Click **Submit** to save the settings.

Step 13 From the left-panel menu, click **Domain Management**. The Domain Management Table page is displayed.

To add domains, see the "Creating, Editing, and Deleting Domains" section on page 6-4.

To view the setting for the domain, click the **View** (eyeglasses) icon next to the domain.

Step 14 Click the **Assign** icon next to each domain name you want to assign to the user account.

To remove the domain from the user account, click the **Assign** icon again.

To assign all domains, click the **Assign All** icon in the task bar. To unassign all domains, click the **Remove All** icon in the task bar.

Step 15 Click **Submit** to save the settings.

To delete a user, in the User Table page, click the **Edit** icon next to the username, and from the User Account page, click the **Trash** icon in the task bar.

Note Deleting a user account from the CLI does *not* delete the corresponding account in the CDSM database. User accounts created in the CDSM should always be deleted from within the CDSM.

Creating, Editing, and Deleting Roles

Although the CDSM provides many types of services, not all users have access to all services. Users are assigned a role, which indicates the services to which they have access. A *role* is a set of enabled services.

Each user account can be assigned zero or more roles. Roles are not inherited or embedded. The CDSM provides one predefined role, known as the *admin role*. The admin role has access to all services and all CDS network entities.

Note The admin user account, by default, is assigned to the role that allows access to all domains and all entities in the system. It is not possible to change the role for this user account.

To create or edit a role, follow these steps:

Step 1 Choose **System > AAA > Roles**. The Roles Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Role page is displayed.

To edit a role, click the **Edit** icon next to the role name.

Step 3 In the **Name** field, enter the name of the role.

Step 4 To expand a listing of services under a category, click the folder, and then check the check box next to the service or services you want to enable for this role. To choose all the services under one category simultaneously, check the check box for the top-level folder.

Step 5 In the **Comments** field, enter any comments about this role.

Step 6 Click **Submit** to save the settings.

To delete a role, in the Roles Table page click the **Edit** icon next to the role name. Once the Role page is displayed, click the **Trash** icon in the task bar.

Creating, Editing, and Deleting Domains

A *domain* is a set of CDS network entities or objects that make up the CDS network. Whereas a role defines which services a user can perform in the CDS network, a domain defines the entities to which the user has access. An *entity* can be a Service Engine, a device group, or a delivery service. These predefined entities are treated like services and can be enabled or disabled when you set up user roles.

When you configure a domain, you can choose to include Service Engines, device groups, or delivery services in the domain.

To create or edit a domain, follow these steps:

Step 1 Choose **System > AAA > Domains**. The Domains Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Domain page is displayed.

To edit a domain, click the **Edit** icon next to the domain name.

Step 3 In the **Name** field, enter the name of the domain.

Step 4 From the **Entity Type** drop-down list, choose Service Engines, Device Groups, or Delivery Services.

Step 5 In the **Comments** field, enter any comments about this domain.

Step 6 Click **Submit** to save the settings. If the entity type you chose has not already been assigned to the domain, then a message displays indicating that the entity type has not been assigned.

Step 7 From the left-panel menu, click **Entity Management**. The Entity Management page is displayed.

Step 8 Click the **Assign** icon (blue cross mark) next to each entity name you want to include. A green arrow wrapped around the blue cross mark indicates an entity will be assigned.

To assign all entities in the domain, click the **Assign All** icon in the task bar.

To remove an entity from the domain, click the **Assign** icon again.

To remove all entities from the domain, click the **Remove All** icon in the task bar.

Step 9 Click **Submit** to save the settings.

To delete a domain, in the Domain Table page click the **Edit** icon next to the domain name. Once the Domain page is displayed, click the **Trash** icon in the task bar.

Changing a Password

If you are a user *without* admin privileges and you are logged in to the CDSM, you can change your own CDSM and CLI user password if you meet the following requirements:

- Your CLI user account and password were created in the CDSM and not in the CLI.
- You are authorized to access the Password page.

Caution We do not recommend changing the CLI user password from the CLI. Any changes to CLI user passwords from the CLI are *not* updated in the management database and are not propagated to the standby CDSM. Therefore, passwords in the management database do not match a new password configured in the CLI.
The advantage of initially setting passwords from the CDSM is that both the primary and the standby CDSMs will be synchronized, and CDSM users will not have to access the CLI to change their passwords.

To change the CDSM and CLI user password for the user account that is currently logged in to the CDSM, do the following:

Step 1 Choose **System > Password**. The Password page is displayed.

Step 2 In the **New Password** field, enter the changed password.

Step 3 In the **Confirm New Password** field, re-enter the password for confirmation.

Step 4 Click **Submit** to save the settings.

Configuring System Settings

This section covers the following topics:

- System Properties
- Configuring Fast SE Offline Detection
- Configuring Distribution QoS
- Configuring Service Routing

System Properties

To modify the system properties, do the following:

Step 1 Choose **System > Configuration > System Properties**. The System Properties page is displayed.

Step 2 Click the **Edit** icon next to the system property you want to change. The Modify Config Property page is displayed.

Step 3 For true or false values, choose a setting from the **Value** drop-down list. For other values, enter a new value. The range is displayed for each numeric value.

Table 6-1 describes the system properties.

Table 6-1 System Properties Fields

Field	Description
cdsm.session.timeout	Length of a Content Distribution Manager session (in minutes).
DeviceGroup.overlap	SE feature overlapping (enable or disable).
System.CmsUnsProgram Sync.Interval	Interval by which CMS synchronizes program import UNS objects (in minutes). The default is 1440 minutes.
System.datafeed.pollRate	Poll rate between the SE or the SR and the CDSM (in seconds).
System.device.recovery.key	Device identity recovery key. This property enables a device to be replaced by another node in the CDS network.
System.healthmonitor.collect Rate	Sets the collect and send rate in seconds for the CMS device health (or status) monitor. If the rate is set to 0, the health monitor is disabled.
System.Icm.enable	Local and CDSM feature (enable or disable). This property allows settings that are configured using the local device CLI or the CDSM to be stored as part of the CDS network configuration data.
System.monitoring.collect Rate	Rate at which the SE collects and sends the monitoring report to the CDSM (in seconds). The default is 300 seconds.
System.monitoring.daily ConsolidationHour	Hour at which the CDSM consolidates hourly and daily monitoring records.

Table 6-1 System Properties Fields (continued)

Field	Description
System.monitoring.enable	SE statistics monitoring (enable or disable).
System.monitoring.monthly ConsolidationFrequency	Frequency (in days) with which the CDSM consolidates daily monitoring records into monthly records.
System.monitoring.record LimitDays	Maximum number of days of monitoring data to maintain in the system.
System.repstatus.update Enabled	Replication status periodic calculations on an SE (enable or disable).
System.repstatus.updateRate	Rate of replication status periodic updates calculated on an SE (in minutes).
System.repstatus.updateRate Sec	Rate of replication status periodic updates calculated on an SE (in seconds). The default is 600 seconds. Setting this rate overrides the update rate set in minutes.
System.repstatus.updateSync Enabled	Sending summary replication status with requested detailed status (enable or disable).

Step 4 Click **Submit** to save the settings.

Configuring Fast SE Offline Detection

You can detect offline SEs more quickly if you enable the Fast Detection of Offline SEs feature.

If Fast Detection of Offline SEs is not enabled, the CDSM waits for at least two “getUpdate request” polling periods before declaring the SE offline.

If Fast Detection of Offline SEs is enabled, the CDSM waits until the value displayed in the Maximum Offline Detection Time field, located on the Configure Fast SE Offline Detection page, is exceeded.

Communication between the SE and CDSM using User Datagram Protocol (UDP) allows faster detection of SEs that have gone offline. UDP heartbeat packets are sent at a specified interval from each SE to the primary CDSM in a CDS network. The primary CDSM tracks the last time it received a UDP heartbeat packet from each SE. If the CDSM has not received the specified number of UDP packets, it displays the status of the nonresponsive SEs as offline. Because UDP heartbeats require less processing than a getUpdate request, they can be transmitted more frequently, and the CDSM can detect offline SEs much faster.

An SE is declared offline when it has failed to contact the CDSM for a getUpdate request (get configuration poll) for at least two polling periods.

Note In CDS networks with heavy traffic, dropped UDP packets can cause the CDSM to incorrectly report the status of SEs as offline. To avoid this problem, configure a higher value for dropped UDP heartbeat packets.

To configure Fast Detection of Offline SEs, do the following:

Step 1 Choose **System > Configuration > Fast SE Offline Detection**. The Configure Fast SE Offline Detection page is displayed.

Note The Fast Detection of Offline SEs feature is in effect only when the CDSM receives the first UDP heartbeat packet and a getUpdate request from an SE.

Step 2 Check the **Enable** check box to enable the CDSM to detect the offline status of SEs quickly.

Note SEs can have their software upgraded without any errors in the status being displayed in the CDSM if you disable Fast Detection of Offline SEs.

Step 3 In the **Heartbeat Rate** field, specify how often, in seconds, the SEs should transmit a UDP heartbeat packet to the CDSM.

Step 4 In the **Heartbeat Fail Count** field, specify the number of UDP heartbeat packets that can be dropped during transmission from SEs to the CDSM before an SE is declared offline.

Step 5 In the **Heartbeat UDP Port** field, specify the CDSM port number that the SEs use to send UDP heartbeat packets.

The **Maximum Offline Detection Time** field displays the product of the failed heartbeat count and heartbeat rate, where:

Maximum Offline Detection Time = Failed heartbeat count * Heartbeat rate

If Fast Detection of Offline SEs is enabled, the CDSM detects SEs that are in network segments that do not support UDP and uses a getUpdate request (get configuration poll) to detect offline SEs.

If the CDSM does not receive regular keep-alive communication from an SE, the CDSM displays the SE as offline after a time period of 2 * (Heartbeat rate) * (Failed heartbeat count).

Step 6 Click **Submit** to save the settings.

Configuring Distribution QoS

The Distribution QoS settings allow you to configure system-wide QoS priorities for delivery service distribution and metadata replication. The delivery service distribution priority (low, medium, or high) is set on the definition page for each delivery service.

Note When a single URL is associated with more than one delivery service, the content is distributed only one time to all of the Service Engines subscribed to each delivery service. When different QoS settings are configured for different delivery services that contain the same content, the delivery service priority setting determines which QoS settings are applied to the content distribution. The delivery service with the higher priority dictates which QoS settings are used.

To configure system-wide QoS settings, do the following:

Step 1 Choose **System > Configuration > Distribution QoS**. The Distribution QoS page is displayed.

Step 2 Check the **Set QoS for Unicast Data** check box to enable system-wide QoS settings for unicast data.

Step 3 Set the QoS value for each priority by choosing the Differentiated Service Code Point (DSCP) value from the QoS value drop-down list or by entering a decimal value in the corresponding field.

Step 4 Check the **Set QoS Value for Metadata Replication** check box to enable QoS settings for metadata replication.

Step 5 Set the Qo**S value for metadata replication** by choosing the DSCP value from the QoS value drop-down list or by entering a decimal value in the corresponding field.

Step 6 Click **Submit** to save the settings.

Configuring Service Routing

The Service Routing menu options consist of the following:

- Coverage Zone File Registration
- Configuring Global Routing

Coverage Zone File Registration

A coverage zone can be associated with one or more SEs: each SE can have its own unique coverage zone, or SEs can be associated with more than one coverage zone and have overlapping coverage zones. For more information about coverage zones, see the "Coverage Zone File" section on page 1-21.

See Appendix C, "Creating Coverage Zone Files," for information about creating a Coverage Zone file.

The system administrator places a Coverage Zone file where the CDSM or individual devices can access the URL. The administrator then registers the Coverage Zone file URL in the CDSM. Coverage Zone files can be applied globally to the entire CDS network, or locally to a specific SR. If a Coverage Zone file is made global, then it is read and parsed by each SR that does not have a Coverage Zone file assigned. If the coverage zone is specified in an individual SR configuration, it is only applied to that particular SR.

You have the choice of using two types of coverage zones:

- Default coverage zones
- User-defined coverage zones

A default coverage zone consists of all the SEs that reside in the same local network segment, or subnet. The CDSM provides a check box to specify whether the default coverage zone is to be used.

A user-defined coverage zone consists of all the SEs that are specified in a Coverage Zone file. This file defines the network segments to be covered in the routing process. The Coverage Zone file is registered with the CDSM and then applied to an SR for routing definitions.

To apply a custom coverage zone to an SR, you first need to register a Coverage Zone file URL in the CDSM. After you have registered the Coverage Zone file URL with the CDSM, you can apply the Coverage Zone file in one of two ways:

- Globally—Deploy the Coverage Zone file across the entire CDS network
- Locally—Deploy the Coverage Zone file on a specific SR

Note If you apply a Coverage Zone file locally for a device, this file overwrites the global Coverage Zone file for that device.

To register a Coverage Zone file, follow these steps:

Step 1 Choose **System > Configuration > Service Routing > Coverage Zone File Registration**. The Coverage Zone File Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Registering Coverage Zone File page is displayed.

Step 3 Choose a file import method from the **File Import Method** drop-down list:

- **Upload**—The upload method allows you to upload a Coverage Zone file from any location that is accessible from your PC by using the browse feature.
- **Import**—The import method allows you to import the Coverage Zone file from an external HTTP, HTTPS, FTP, or CIFS server.

When you choose a method, the page refreshes and displays the configuration fields that are associated with the method that you chose.

Step 4 Enter the fields as appropriate. Table 6-2 describes the upload method fields. Table 6-3 describes the import method fields.

Table 6-2 Upload Method for Coverage Zone Files

Property	Description
Coverage Zone File Upload	Local directory path to the Coverage Zone file. To locate the file, use the **Browse** button.
Destination Filename	Name of the Coverage Zone file. This field is filled in automatically with the filename from the local directory path.

Table 6-3 Import Method for Coverage Zone Files

Property	Description
Coverage Zone File URL	The URL where the Coverage Zone file is located, including path and filename.
Destination File Name	Name of the Coverage Zone file.
Update Interval (minutes)	Frequency with which the CDSM looks for changes to the Coverage Zone file. The default value is 10 minutes.
Username	Name of the user to be authenticated when fetching the Coverage Zone file.
Password	User password for fetching the Coverage Zone file.

Table 6-3 Import Method for Coverage Zone Files

Property	Description
NTLM user Domain	NT LAN Manager (NTLM) user domain name for NTLM authentication.
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method, and the username and password information can be passed to the origin server in clear text with a basic authentication header.

Step 5 To save the settings, click **Submit**.

Configuring Global Routing

After you have registered the Coverage Zone file, you can use this file as your global routing configuration.

To set a global Coverage Zone file, do the following:

Step 1 Choose **System > Configuration > Service Routing > Global Routing Config**. The Set Global Coverage Zone File page is displayed.

Step 2 From the **Coverage Zone File** drop-down list, choose a Coverage Zone file.

Step 3 In the **Keep Alive Port** field, enter a port number that is used by the CDS network to listen in on routing requests.

Step 4 In the **DNS TTL** field, configure the time period (in seconds) for caching DNS replies. Enter a number from 0 to 2147483647. The default is 60 seconds.

Step 5 Click **Submit** to save settings.

To apply a Coverage Zone file to an individual SR for local coverage zone configuration, see the "Configuring the Service Router" section on page 4-87.

CHAPTER 7

Monitoring the Internet Streamer CDS

The CDSM provides tools that can be used for system monitoring and system diagnostics. The topics covered in this chapter include:

- System Monitoring, page 7-1
- Device Monitoring, page 7-6
- Reports, page 7-14
- Delivery Service Monitoring, page 7-17
- Viewing Statistics, page 7-27
- Transaction Logs, page 7-31

System Monitoring

System monitoring consists of the following:

- System Status
- System Home Page
- System Audit Logs

System Status

The CDSM displays the system status in the System Status bar that is located above the navigation tabs in every window. The System Status bar presents the overall device and content health of the system. You can use this feature to monitor devices and content replication in your CDS network. The System Status bar helps you immediately identify any problems on the network, allowing you to act and respond to problems quickly.

The system status reporting mechanism uses four alarm lights to identify problems that need to be resolved. Each light represents a different alarm level, as follows:

- Green—No alarms (the system is in excellent health)
- Yellow—Minor alarms
- Orange—Major alarms
- Red—Critical alarms

When you roll your mouse over an alarm light in the System Status bar, a pop-up message provides further details about the device or delivery service status. See Figure 7-1.

Figure 7-1 System Status Bar

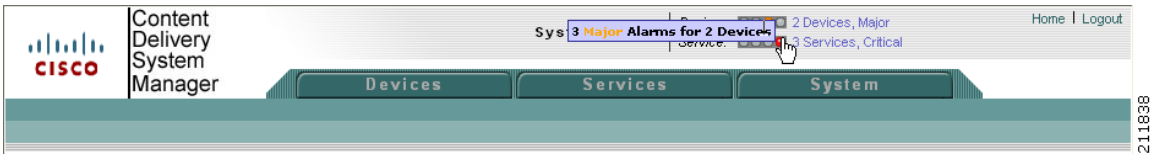

When you click the alarm light, a troubleshooting window opens (Troubleshooting Devices or Troubleshooting Service), listing the individual devices or delivery services that need attention.

Figure 7-2 Troubleshooting Tools Menu

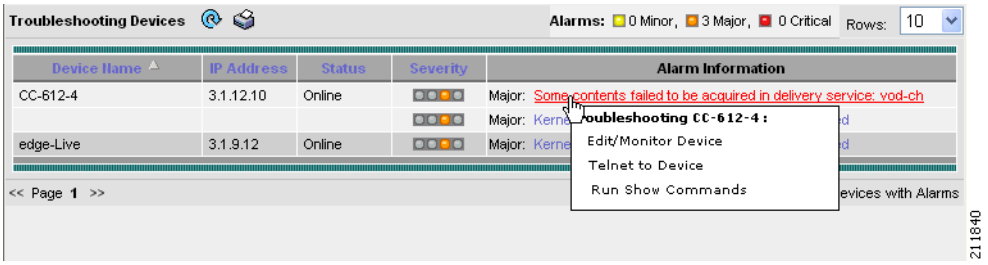

When you roll your mouse over an item under the Alarm Information column in the Troubleshooting Devices or Troubleshooting Services window, the Troubleshooting Tools menu is displayed. The Troubleshooting Tools menu provides links to all the diagnostic tools, troubleshooting tools, and monitoring applications for troubleshooting and resolving the problem. Figure 7-2 shows the Troubleshooting Tools menu for device alarms.

Device Alarms

Device alarms are associated with device objects and pertain to applications and services running on SEs, SRs, and CDSMs. Device alarms are defined by the reporting application or service. Device alarms can also reflect reporting problems between the device and the CDSM. (See Table 7-1.)

Table 7-1 Device Alarms for Reporting Problems

Alarm	Alarm Severity	Device Status	Description
Device is offline	Critical	Offline	The device has failed to communicate with the CDSM.
Device is pending	Major	Pending	The device status cannot be determined.
Device is inactive	Minor	Inactive	The device has not yet been activated or accepted by the CDSM.
Device has lower software version	Minor	Online	The device is not interoperable with the CDSM because it has an earlier software version.

Troubleshooting Devices Using the System Status Bar

To troubleshoot a device from the System Status bar, do the following:

Step 1 In the **System Status** bar, click the **Devices** alarm light or click the **Device** link. The Troubleshooting Devices window is displayed.

Step 2 In the Alarm Information column, hold your mouse over the alarm message until the Troubleshooting Tools menu is displayed. See Figure 7-2.

Step 3 Click the troubleshooting tool you want to use. The link takes you to the corresponding page in the CDSM. Table 7-2 describes the tools available for all device alarms.

Table 7-2 Troubleshooting Tools for Device Alarms

Item	Navigation	Description
Edit/Monitor Device	Device home page	Displays device home page
Telnet to Device	Opens an Telnet window	Initiates an Telnet session using the device IP address
Run Show Commands	**Devices > Device Monitoring > Show/Clear Commands > Show Commands**	Displays device show command tool

Service Alarms

Service alarms pertain to content replication problems and are associated with delivery services. Service alarms are raised by the CDSM based on replication status reports, or by the SE health manager based on acquisition and distribution errors.

If the same fault is reported by the replication status and by the SE health manager, the CDSM reports both; one appears as the true alarm and the other as an error. The CDSM does not correlate nor attempt to consolidate the errors generated by the replication status and by the SE health manager.

To troubleshoot service replication issues from the System Status bar, do the following:

Step 1 In the System Status bar, click the **Services** alarm light or click the **Service** link. The Troubleshooting Services window is displayed. Table 7-3 lists the service alarms.

Table 7-3 Service Alarms for Delivery Service Replication Status

Alarm	Severity	Description
Replication Status is Failed	Critical	The number of SEs in the delivery service that failed to replicate the content is greater than zero.
Replication Status is Pending	Minor	The number of SEs in the delivery service with content replication status unknown is greater than zero.

Step 2 In the Alarm Information column, hold your mouse over the alarm message until the Troubleshooting Tools menu is displayed.

Step 3 Click the troubleshooting tool that you want to use. The link takes you to the corresponding page in the CDSM. Table 7-4 describes the tools available for all service alarms.

Table 7-4 Troubleshooting Tools for Content Alarms

Item	Navigation	Description
View Replication Status	**Services > Delivery Services > Replication Status**	Displays second-level replication status for a delivery service.
Edit Delivery Service	**Services > Delivery Services> Definition**	Opens the Delivery Service Definition page.

System Home Page

The System Home page (Figure 7-3) provides overall system performance graphs, and overall system information on configuration and software versions running on the CDS devices. Clicking the links for devices, delivery services, and programs take you to the corresponding table pages.

Figure 7-3 System Home Page

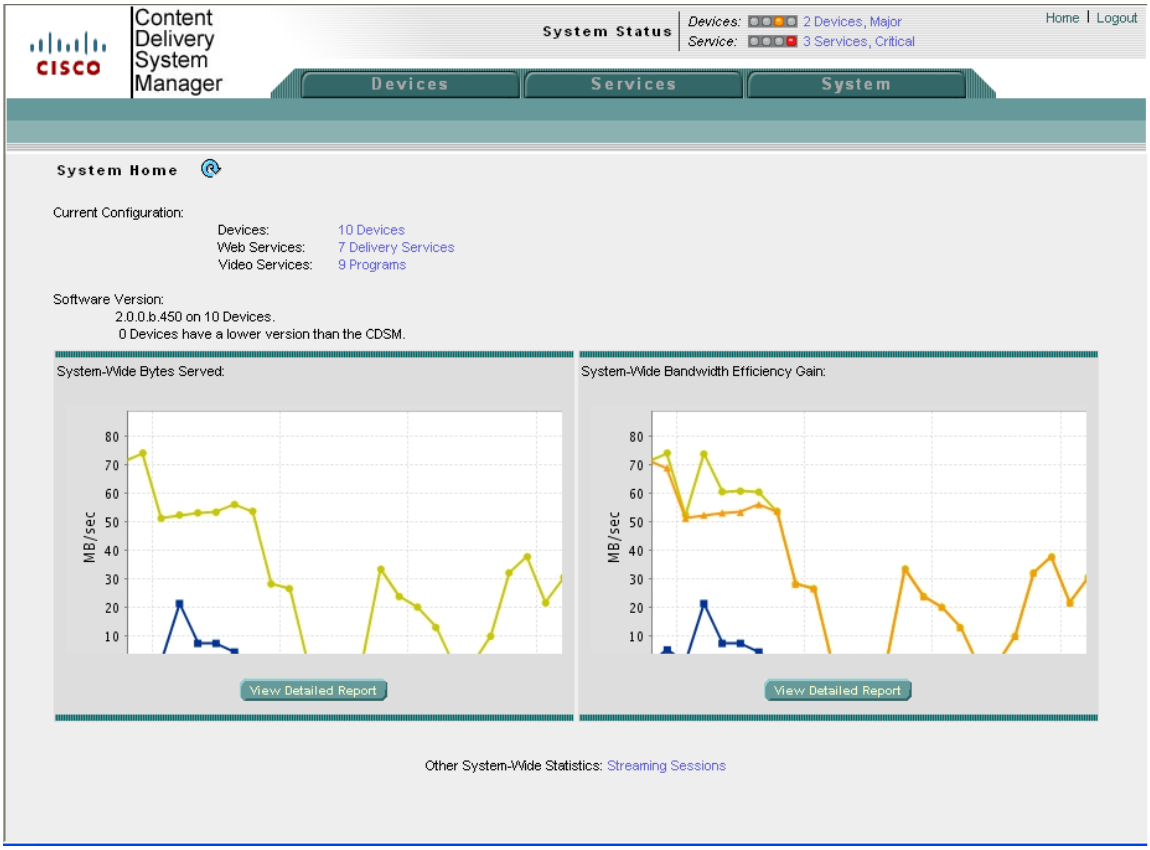

The information displayed in the graphs is based on a snapshot of your CDS network and represents the state of your SEs at the end of every two polling periods. You can change the interval between polls by changing the System.datafeed.pollRate field in **System > Configuration > System Properties**. The default polling rate is 300 seconds (5 minutes).

To change the report settings for the System-Wide Bytes Served or System-Wide Bandwidth Efficiency Gain graphs, click **View Detailed Report**. Clicking the **Streaming Sessions** link at the bottom of the home page opens the System-Wide Streaming Sessions Report page (Figure 7-4). For more information about these reports, see the "Reports" section on page 7-14.

Figure 7-4 System-Wide Streaming Sessions Report Page

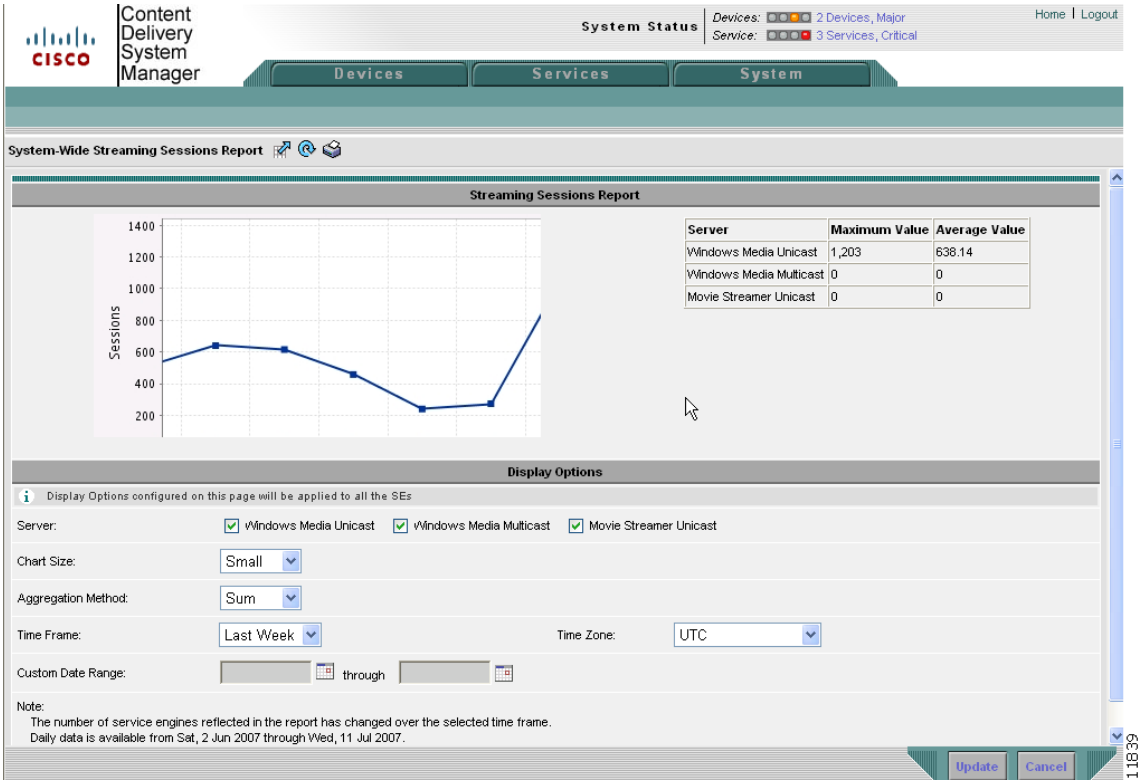

System Audit Logs

The CDSM logs user activity in the system. The only activities that are logged are those that change the CDS network. This feature provides accountability for users actions (for example, which user did what and when). Logged activities include the following:

- Creation of CDS network entities
- Modification and deletion of CDS network entities
- System configurations

To view audit trail logs, do the following:

Step 1 Choose **System > Logs > Audit Trail Logs**. The Audit Log page is displayed. All logged transactions in the CDSM are listed by date and time, user, actual transaction that was logged, and the IP address of the machine that was used.

Step 2 To determine the number of rows that you want to display, choose a number from the Rows drop-down list.

Device Monitoring

This section covers the following topics:

- Devices Table
- Device Home Page
- Using show and clear Commands
- CPU Utilization

For more detailed statistics on HTTP, Web Media, Movie Streamer, and Flash Media Streaming traffic, see the "Viewing Statistics" section on page 7-27.

Devices Table

The Devices Table page displays all devices registered in the CDS network (Figure 7-5).

Figure 7-5 Devices Table Page

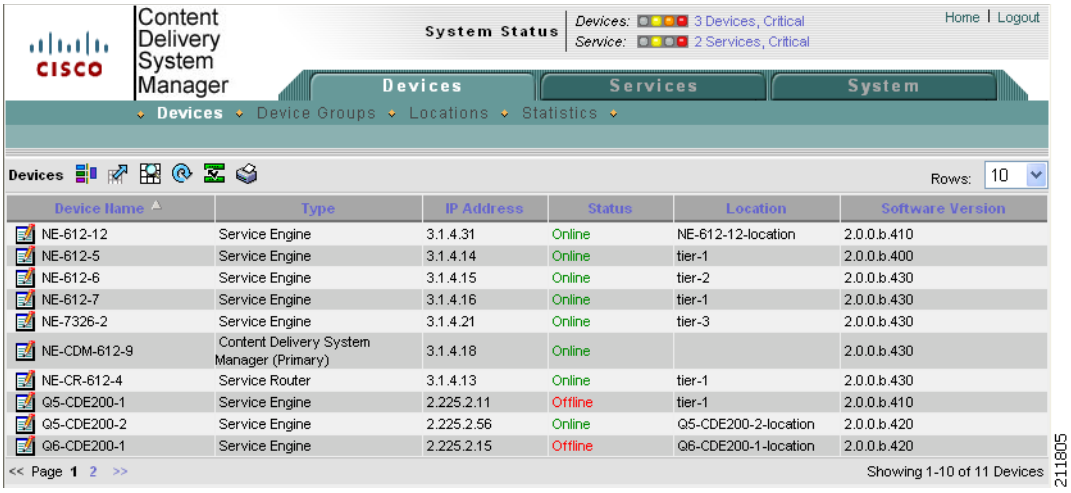

Table 7-5 describes the Device Table columns. You can sort the information in the table by clicking on any column title. The table can be sorted in ascending or descending order for each column.

The table defaults to listing ten rows. You can change the number of rows by clicking the Rows drop-down list. The bottom of the table lists the page number and the total number of pages, as well as how many items are showing out of the total number of items.

Table 7-5 Device Table Columns

Column Heading	Description
Device Name	Host name of the device.
Type	Device type: SE, SR, CDSM (Primary), CDSM (Secondary)
IP Address	Primary IP address of the device.
Status	Status is one of the following: • Online—Device has been activated through the CDSM and is able to send and receive data and control traffic. • Offline—Device has failed to communicate with the CDSM. • Pending—Device status cannot be determined. The device could be in the process of being activated by the CDSM • Offloading—Device is in the Server Offload state. See the Server Offload field in Table 4-2 on page 4-8 for more information. To monitor the current streams on an SE during the Server Offload state, view the statistics for each protocol engine (for example, Movie Streamer), specifically the fields noted in Table 7-22 on page 7-27. Once all protocol engines have finished streaming, you can perform maintenance or upgrade the software on the device. For information about upgrading the software, see the "Upgrading the Software" section on page 8-6.
Location	Location the device is assigned to.
Software Version	Device software version.

The task bar options provide other table manipulations, including filtering, exporting the table, refreshing the table, viewing all items, activating all inactive SEs, and printing. The filter option allows you to filter on device name, device type, or status. To display all items after a filter has been set, click the **View All** icon.

Device Home Page

The Device home page (Figure 7-6) provides alarm status and information on the device. Only basic information is displayed for the SR and CDSM.

Clicking the Delivery Services and Device Groups links on the home page for an SE takes you to a corresponding table listing all the delivery services or device groups in the CDS, and which ones the SE is assigned to. Through this page, you can assign the device to additional delivery services or device groups by clicking the icon next to the applicable delivery services or device groups and submitting your selection.

You can update the device software, and telnet to the device from the Device home page. For more information about updating the software, see the "Software Upgrade" section on page 8-1.

Figure 7-6 SE Device Home Page

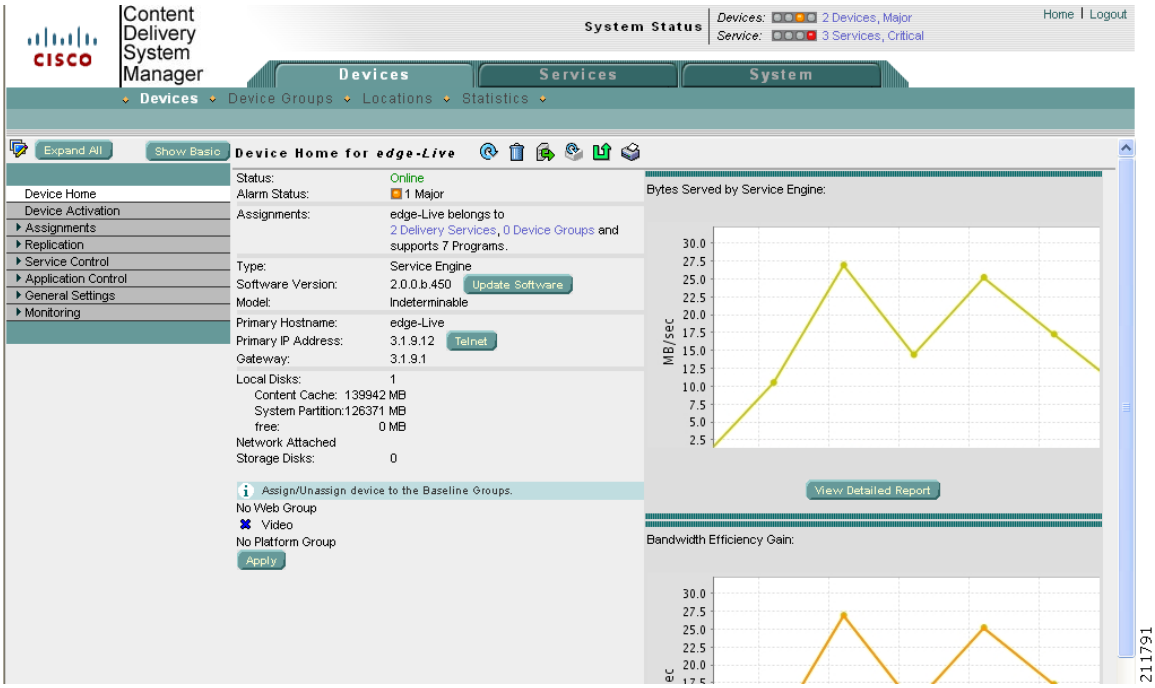

The Bytes Served by Service Engine and the Bandwidth Efficiency Gain graphs are also displayed. For more information, see the "Reports" section on page 7-14.

The Device home page for an SE or an SR provides several task bar options. Table 7-6 describes these options. The CDSM home page has a subset of the task bar options.

Table 7-6 SE Device Home Page Task Bar Icons

Rollover Label	Description
Refresh	Refreshes the page.
Delete Device	Deletes the device. See the "Deleting a Device" section on page 8-10.
Update Application Statistics	The device statistics are updated at a configurable time interval, which is set in the System Configuration page (System.monitoring.collectRate). See the "Configuring System Settings" section on page 6-6. To see the latest statistics immediately, without waiting for the time interval to elapse, click this icon.
Force Full Database Update	Forces a full database update from CDSM to the SE. If the CDSM and SE databases are not synchronized as a result of network errors or other errors, you can synchronize them by clicking this icon.
Reload Device	Reboots the device. See the "Rebooting Devices" section on page 8-9.
Print	Prints the home page.

Using show and clear Commands

The **show** and **clear** commands offer more detailed monitoring of the device. Table 7-7 lists only the **show** command parameters where arguments are required or are optional. Table 7-8 lists only the **clear** command parameters where arguments are required. A full list of the **show** and **clear** commands is available from the drop-down list on the respective page.

Using the CDSM show or clear Command Tool

To use the CDSM show or clear command tool, follow these steps:

Step 1 Choose **Devices > Devices > Monitoring > Show/Clear Commands** and then click either **Show Commands** or **Clear Commands**.

Step 2 From the drop-down list, choose a command.

Step 3 Enter arguments for the command, if any.

Step 4 Click **Submit** to display the command output. The results of the command are displayed in a new window.

Table 7-7 show Command Arguments

show Command	Arguments	Device
access-list	**300**	SE
acquirer	[**delivery-service** [**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*] | **progress** [**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*] | **proxy authentication**]	SE
alarms	[**critical** | **detail** | **history** | **major** | **minor** | **status**]	SE, SR, CDSM
authentication	**http-request** | **user**	SE, SR, CDSM
bandwidth	[**flash-media-streaming** | **movie-streamer** | **wmt**]	SE
bitrate	[**wmt**]	SE
cache	[**content** *1-1000*]	SE
cdn-statistics	{**flash-media-streaming** {**device-group-name** *device-group-name* | **device-groups** | **service-engines**} **movie-streamer** {**service-engines** | **device-group-name** *groupname* | **device-groups**} | **http** {**service-engines** | **device-group-name** *groupname* | **device-groups**} | **wmt** {**service-engines** | **device-group-name** *groupname* | **device-groups**}}	CDSM
cdnfs	**usage** | **volumes**	SE
clock	[**detail** | **standard-timezones** {**all** | **details** *timezone* | **regions** | **zones** *region-name*}]	SE, SR, CDSM

Table 7-7 show Command Arguments (continued)

show Command	Arguments	Device
cms	{**database** {**content** {**dump** *filename* | **text** | **xml**} | **maintenance** [**detail**]} | **info** | **processes**}	SE, SR, CDSM
content	**all** | **url**	SE
device-mode	{**configured** | **current**}	SE, SR, CDSM
disks	[**current** | **details** | **failed-sectors** [*disk_name*] | **raid-state** | **SMART-info** [**details**]]	SE, SR, CDSM
distribution	[**delivery-services** [**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*]] [**forwarder-list** [**delivery-service-id** *delivery-service-num* [**detail**] | **delivery-service-name** *delivery-service-name* [**detail**] | **detail**]] [**location** {**forwarder-load-weight** | **live-load-weight** | **location-leader-preference**} [**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*]] **[object-status** *object-url*] [**processes**] **[remote** *ip-address* {**metadata-sender delivery-service-id** *delivery-service-num* [**start-generation-id** *gen-id* **end-generation-id** *gen-id*] | **unicast-sender delivery-service-id** *delivery-service-num* {**cdn-url** *cdn_url* | **probe** | **relative-cdn-url** *cdn_url*}}] [**remote traceroute** {**forwarder-next-hop delivery-service-id** *delivery-service-num* {**max-hop** *maxhop_num* | **trace-till-good** | **trace-till-root**} | **unicast-sender delivery-service-id** *delivery-service-num* {**cdn-url** *cdn-url* | **probe** | **relative-cdn-url** *cdn-url*} {**max-hop** *maxhop_num* | **trace-till-good** | **trace-till-root**}}]	SE, SR
flash-streaming-media	[**livestreams** [**detail** {*filename*}] | **mtrack**]	SE, SR
http	{**age-mult** | **all** | **cache-cookie** | **cache-on-abort** | **object** | **proxy** | **reval-each-request** | **ttl**}	SE, SR
icap	[**service** *service_name*]	SE
interface	{**FastEthernet** *slot/port* | **FibreChannel** *slot/port* | **GigabitEthernet** *slot/port* | **ide** *control_num* | **PortChannel** *port-num* | **scsi** *device_num* | **Standby** *group_num* | **usb**}	SE, SR, CDSM
ip	**access-list** [*acl-name* | *acl-num*] **routes**	SE, SR, CDSM
movie-streamer	[**bandwidth** | **cache** | **proxy**]	SE
network-filesystem	{**client** {**all** | **cifs** | **nfs** | **samba**} | **server** {**nfs**}}	SE
ntp	**status**	SE, SR, CDSM

Table 7-7 show Command Arguments (continued)

show Command	Arguments	Device
processes	[**cpu** | **debug** *pid* | **memory** | **system** [**delay** *1-60* | **count** *1-100*]]	SE, SR, CDSM
proxy-protocols	**outgoing-proxy**	SE
rea	**agent**	SE
rtsp	{**gateway**}	SE
rule	{**action** {*action-type* | **all**} [**protocol** {**http** | **rtsp**}] | **all** | **pattern-list** {*1-512 pattern-type* | **all**}}	SE
service-router	On SE: {**keepalive-interval** | **service-monitor**} On SR: [**content-based-routing** | **dns-ttl** | **forwarding** [**content-origin** {*content-origin*}] | **lastresort** | **load** {**all** | **sename** {*sename*} } | **location-based-routing** | **routes** [**content-origin** {*content-origin*}] | **services** {**all** | **sename** {*sename*} | **summary**]	SE, SR
services	{**ports** [*port-num*] | **summary**}	SE, SR, CDSM
snmp	{**alarm-history** | **engine ID** | **event** | **group** | **stats** | **user**}	SE, SR, CDSM
statistics	**access-lists 300** **acquirer** [**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name* | **contents** {**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*} | **errors** {**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*} | **job-list** {**delivery-service-id** *delivery-service-num* | **delivery-service-name** *delivery-service-name*}] **authentication** **cdnfs** **distribution** {**all** | **errors** {**delivery-service-id** *delivery-service-num* | **delivery-service-name** *name*} | **metadata-receiver** | **metadata-sender** | **unicast-data-receiver** [**delivery-service-id** *delivery-service-num* [**pending-queue** *jobs* | **suspended-queue** *jobs* | **waiting-queue** [**first** | **last**] [*jobs*]] | **delivery-service-name** [**pending-queue** *jobs* | **suspended-queue** *jobs* | **waiting-queue** [**first** | **last**] [*max_jobs*]] | **hot-forwarders** [*forwarder_id* | *forwarder_name*] {**idle-queue** | **priority-queue**} *delivery-services* | **idle-forwarders** *max_idle_forwarders*] | **unicast-data-sender**} **flash-media-streaming** [**connections** | **errors** | **performance** | **requests**] **http** {**ims** | **object** | **pcmm** | **performance** | **requests** | **rule**}	SE, SR, CDSM

Table 7-7 show Command Arguments (continued)

show Command	Arguments	Device
statistics	**icap** **icmp** **ip** **movie-streamer** {**all** | **bw-usage** | **errors** | **performance** | **requests** | **rule**} **netstat** **qos** {**policy**} **radius** **replication** **rtsp server**[1] **services** **service-router** {**all** | **content-origin** {*content-origin*} | **dns** | **history** | **keepalive** | **se** {*se-name*} | **summary**} **snmp** **tcp** **transaction-logs** **udp** **wmt** {**all** | **bytes** [**incoming** | **outgoing**] | **cache** | **errors** | **multicast** | **requests** | **rule** | **savings** | **streamstat** [**incoming** | **live** | **outgoing** | **stream-id** *1-999999*] | **usage**}	SE, SR, CDSM
tech-support	[**list-files** *list-file-directory* **page** | **service** {**acquisition-distribution** | **authentication** | **cms** | **cse** | **dns** | **ftp** | **http** | **icap** | **kernel** | **rules** | **wmt**}]	SE, SR, CDSM
user	{**uid** *number* | **username** *name*}	SE, SR, CDSM
users	{**administrative**}	SE, SR, CDSM
wmt	[**bandwidth** [**incoming bypass-list**] | **broadcast** | **detail** | **diagnostics** {**header-info** {**stream-file** *word*}| **nsc-file** *nsc-filename*} | **http allow extension** | **multicast** | **proxy**]	SE

1. For CDS Releases 2.0 and 2.1, enter the **show statistics rtsp server movie-streamer all** command. For CDS Release 2.2 and 2.3, enter the **show statistics movie-streamer all** command.

Note All WMT playable contents can be delivered by either HTTP or RTSP, based on the request. Any content that is cached by the WMT is stored using the RTSP scheme, regardless of whether the content was cached due to an HTTP or RTSP request. Therefore, in the **show** command, the content displays as RTSP.

Table 7-8 clear Command Arguments

clear Command	Arguments	Device
cache	[**all** | **content** {*1-1000000*}]	SE
content	**url** *url*	SE
ip	**access-list counters** *1-99*	SE, SR, CDSM
statistics	{**access-lists 300** | **all** | **authentication** | **bandwidth advanced errors** | **distribution** {**all** | **metadata-receiver** | **metadata-sender** | **unicast-data-receiver** | **unicast-data-sender**} | **history** | **http** {**all** | **ims** | **object** | **pcmm** | **requests** | **rule**} | **icap** | **icmp** | **ip** | **qos policy-service** | **radius** | **rtsp server**[1] | **movie-streamer** | **rule** {**action** *action-type* | **all** | **pattern** {*1-512* | **all**} | **rtsp**} | **running** | **tcp** | **transaction-logs** | **udp** | **wmt**}	SE, SR, CDSM
users	{**administrative** | **request-authenticated**}	SE, SR, CDSM
wmt	{**stream-id** *1-999999*}	SE

1. For CDS Releases 2.0 and 2.1, enter the **show statistics rtsp server movie-streamer all** command. For CDS Release 2.2 and 2.3, enter the **show statistics movie-streamer all** command.

CPU Utilization

The CPU Utilization report displays the CPU usage for the SE.

To view the CPU Utilization report for an SE, do the following:

Step 1 Choose **Devices > Devices > Monitoring > Statistics > CPU Utilization**. The CPU Utilization Report page is displayed.

Step 2 Enter the settings as appropriate. Table 7-9 describes the report settings.

Table 7-9 CPU Utilization Report Settings

Field	Description
Chart Size	The chart display size choices are small, medium, or large.
Time Frame	The time frame options are last hour, last day, last week, last month, or custom.
Time Zone	The time zone choices are SE local time, CDSM local time, or UTC.
Custom Date Range	The custom date range is used when Time Frame is set to custom. Enter the dates, beginning and end, for the chart in the mm/dd/yyyy format, or choose the dates by using the calendar icons.

Step 3 Click **Update** to see the report.

To export the report to a CSV (comma-separated value) file, click the **Export** icon in the task bar. A dialog box is displayed. Choose either **Open** or **Save**.

If you choose **Open**, the tabular report is displayed in a new browser window.

If you choose **Save**, you are prompted to choose a location where to save the file. The file can be opened with any spreadsheet program.

To print the report, click the **Print** icon in the task bar.

Reports

There are three reports available for monitoring traffic in graphical or tabular format:

- Bytes Served
- Bandwidth Efficiency Gain
- Streaming Sessions

The reports have the following three scopes:

- System-wide
- Location
- Service Engine

To access the system-wide reports, click the **Home** link in the upper-right corner of the CDSM browser window. To change the report parameters for the System-Wide Bytes Served or System-Wide Bandwidth Efficiency Gain graphs, click **View Detailed Report**. Clicking the **Streaming Sessions** link opens the System-Wide Streaming Sessions page.

Note If the report states, "Insufficient data. Please make sure NTP is configured on the SE." Be sure NTP is configured for each device that is contributing data to the report. See the "Configuring NTP" section on page 4-58 for more information.

To access reports covering activity for a location, do the following:

Step 1 Choose **Devices > Locations**. The Location Table page is displayed.

Step 2 Click the **Edit** icon next to the location name. The Location page is displayed.

Step 3 Choose **Statistics** and choose one of the following reports: **Bytes Served**, **Bandwidth Efficiency Gain**, or **Streaming Sessions**.

To access reports covering activity for an SE, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed.

Step 2 Click the **Edit** icon next to the device name. The Device home page is displayed.

Step 3 Choose **Statistics** and choose one of the following reports: **Bytes Served**, **Bandwidth Efficiency Gain**, or **Streaming Sessions**.

To export the report to a CSV (comma-separated value) file, click the **Export** icon in the task bar. A dialog box is displayed. Choose either **Open** or **Save**.

If you choose **Open**, depending on your browser, the tabular report is displayed in either a new browser window or the same browser window.

If you choose **Save**, you are prompted to choose a location where to save the file. The file can be opened with any spreadsheet program.

To print the report, click the **Print** icon in the task bar.

The reports are described in the following sections.

Bytes Served

The Bytes Served report provides information about the number of bytes per second delivered to client devices over a specified period of time. To change the report settings and view the changes, navigate to the page using the instructions provided at the beginning of the "Reports" section on page 7-14.

Table 7-10 describes the report settings.

Table 7-10 Bytes Served Report Settings

Field	Description
Server	The options are HTTP, Windows Media, Movie Streamer, or Flash Media Streaming. Check the check boxes next to the protocol engines you want to include in the graph.
Chart Style	The options are line or area.
Chart Size	The chart display size choices are small, medium, or large.
Aggregation Method	For system-wide and location reports only. Choices are sum or average, where sum gives you the sum total of all bytes served in the system or location, and average divides the sum total by the number of SEs in the system or location.
Include Child Location	For location report only. If checked, all child locations are included in the report.
Time Frame	The time frame options are last hour, last day, last week, last month, or custom.
Time Zone	The time zone choices are SE local time, CDSM local time, or UTC.
Custom Date Range	The custom date range is used when Time Frame is set to custom. Enter the dates, beginning and end, for the chart in the mm/dd/yyyy format, or choose the dates by using the calendar icons.

Tip Set the Chart Style to medium to see the legend and timeline across the bottom.

Bandwidth Efficiency Gain

After an SE has been in use for some time and has collected statistics, the Bandwidth Efficiency Gain report can demonstrate the value of the SE in terms of bandwidth savings.

Table 7-11 describes the report settings.

Table 7-11 Bandwidth Efficiency Gain Report Settings

Field	Description
Series	The series options are In, Out, and Efficiency Gain. The In option creates a graph for bandwidth used for incoming data. The Out option is for outgoing data, and Efficiency Gain is the combination of the two.
Chart Size	The chart display size choices are small, medium, or large.
Aggregation Method	For system-wide and location reports only. Choices are sum or average, where sum gives you the sum total of all bytes served in the system or location, and average divides the sum total by the number of SEs in the system or location.
Include Child Location	For location report only. If checked, all child locations are included in the report.
Time Frame	The time frame options are last hour, last day, last week, last month, or custom.
Time Zone	The time zone choices are SE local time, CDSM local time, or UTC.
Custom Date Range	The custom date range is used when Time Frame is set to custom. Enter the dates, beginning and end, for the chart in the mm/dd/yyyy format, or choose the dates by using the calendar icons.

Tip Set the Chart Size to medium to see the legend and timeline across the bottom.

Streaming Sessions

The Streaming Sessions report lists the total number of streaming sessions in progress at the collection time. It allows you to plan for future hardware provisioning and licensing requirements based on utilization data. Table 7-12 describes the report settings.

Table 7-12 Streaming Sessions Report Settings

Field	Description
Server	The options are Windows Media unicast, Windows Media multicast, Movie Streamer unicast, or Flash Media unicast. Check the check boxes next to the streaming types you want to include in the graph.
Chart Size	The chart display size choices are small, medium, or large.
Aggregation Method	For system-wide and location reports only. Choices are sum or average, where sum gives you the sum total of all bytes served in the system or location, and average divides the sum total by the number of SEs in the system or location.
Include Child Location	For location report only. If checked, all child locations are included in the report.
Time Frame	The time frame options are last hour, last day, last week, last month, or custom.

Table 7-12 Streaming Sessions Report Settings (continued)

Field	Description
Time Zone	The time zone choices are SE local time, CDSM local time, or UTC.
Custom Date Range	The custom date range is used when Time Frame is set to custom. Enter the dates, beginning and end, for the chart in the mm/dd/yyyy format, or choose the dates by using the calendar icons.

Note Streaming Sessions statistics report for Movie Streamer is only available for unicast. When a client is joining a multicast group for multicast streaming, CDS Movie Streamer only knows that a client is downloading the SDP file, but no information is exchanged between the client and Movie Streamer on the streaming data session; therefore there are no session statistics for multicast Movie Streamer sessions.

Tip Set the Chart Size to medium to see the legend and timeline across the bottom.

Delivery Service Monitoring

This section covers the following topics:

- Delivery Services Table
- Replication Status for a Delivery Service

Delivery Services Table

The Delivery Services Table page lists all delivery services on the system and displays the replication status information for each delivery service. This display summarizes the replication status of all SEs associated with a specific delivery service in a given state.

To view system-wide replication status for each delivery service, do the following:

Step 1 Choose **Services > Delivery Services** to display the Delivery Services Table page. See Figure 7-7.

Figure 7-7 Delivery Services Table Page

Step 2 View the replication status information for each delivery service. Table 7-13 describes the status information that is displayed on this page.

Table 7-13 System-Wide Replication Status by Delivery Service

Column Heading	Description
Delivery Service	Name of the delivery service.
Type	Type of delivery service. The delivery service types are Live and Content.
Content Origin	Name of the Content Origin assigned to the delivery service.
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: • Green—No errors encountered. • Yellow—Only minor errors encountered. • Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive SE. For details of the errors, click the status light for a particular delivery service, which takes you to the Replication Status for Delivery Service page. (See Table 7-14 for a description of status errors and their corresponding status lights.)

Table 7-13 System-Wide Replication Status by Delivery Service (continued)

Column Heading	Description
State	State of the delivery service. States are reported for the Content Acquirer and for receiver SEs. (See Table 7-15 for a definition of the different delivery service states.) The state is also a link to the Replication Status for Delivery Service page that provides a more detailed view of the replication status for the delivery service. (See Figure 7-8.)
Manifest State	State of the Manifest file. States reported are as follows: • Fetching—The Manifest file is being fetched. • Fail Fetching—The Manifest file has failed to be fetched. • Parsing—The Manifest file is being parsed. • Fail Parsing—The Manifest file has failed to be parsed. • Completed—The Manifest file was successfully fetched and parsed. • No Status Reported—Content Acquirer is in a Pending or Disabled state.

Table 7-14 describes the status errors and their corresponding status lights.

Table 7-14 Delivery Service Status Errors

Status Light	Error	Description
Yellow	Manifest retrieval error	The Content Acquirer cannot retrieve the Manifest file for one or two consecutive attempts.
Red	Manifest retrieval error	The Content Acquirer cannot retrieve the Manifest file for three consecutive attempts.
Red	Manifest syntax error	The Content Acquirer fails to parse the Manifest file.
Red	Crawl job processing error	The Content Acquirer encounters problems while crawling for content.
Red	Acquisition or content replication error	The SE fails to obtain the content.
Red	Disk quota exceeded error	The SE cannot store or process the content because there is no more disk space available.
Yellow	Replication status update error	Content replication failed for one or two consecutive attempts.
Red	Replication status update error	Content replication failed for three or more consecutive attempts.
Red	SE unreachable error	The SE is offline or the SE has not responded to replication status requests for three consecutive polling periods.
Red	Root SE failover	The Content Acquirer has failed over to a temporary Content Acquirer. Receiver SEs have not identified a valid Content Acquirer.
Red	Receiver SE device or delivery service error	Receiver SE is not reporting replication status or any other content replication problem.

Table 7-15 defines the different delivery service states.

Table 7-15 Delivery Service States in Replication Status

State	Description
Completed	All receiver SEs are in the Completed state, and the Content Acquirer is in the Completed, Re-checking Content, Retrieving Manifest, or Processing Manifest state. (See Table 7-18 on page 7-22 for a description of SE states.) When the Content Acquirer in the Re-checking Content state determines that new content needs to be acquired, the delivery service state changes to In Process.
In Process	In Process can mean: • The Content Acquirer is in the Retrieving Manifest, Processing Manifest, Acquiring Content, or Re-checking Content state. • Any receiver SE is in the Pending Update from Content Acquirer, Replicating, or Recovering from Failure state. • The Content Acquirer has failed and receiver SEs are still reporting status.
Failed	Failed can mean: • An acquisition or content replication error has occurred. (See Table 7-14 on page 7-19.) • An SE has gone offline or has not reported status in three consecutive polling periods. • The delivery service has more than one Content Acquirer • The delivery service has no Content Acquirer, but has receiver SEs reporting replication status.

Replication Status for a Delivery Service

To view the replication status for a delivery service, you can either click the alarm light or **Replication Status** link in the Delivery Services Table, or click the **Replication Status** option from the Delivery Service left-panel menu. Figure 7-8 shows the Replication Status page for a delivery service.

Figure 7-8 Delivery Service Replication Status Page

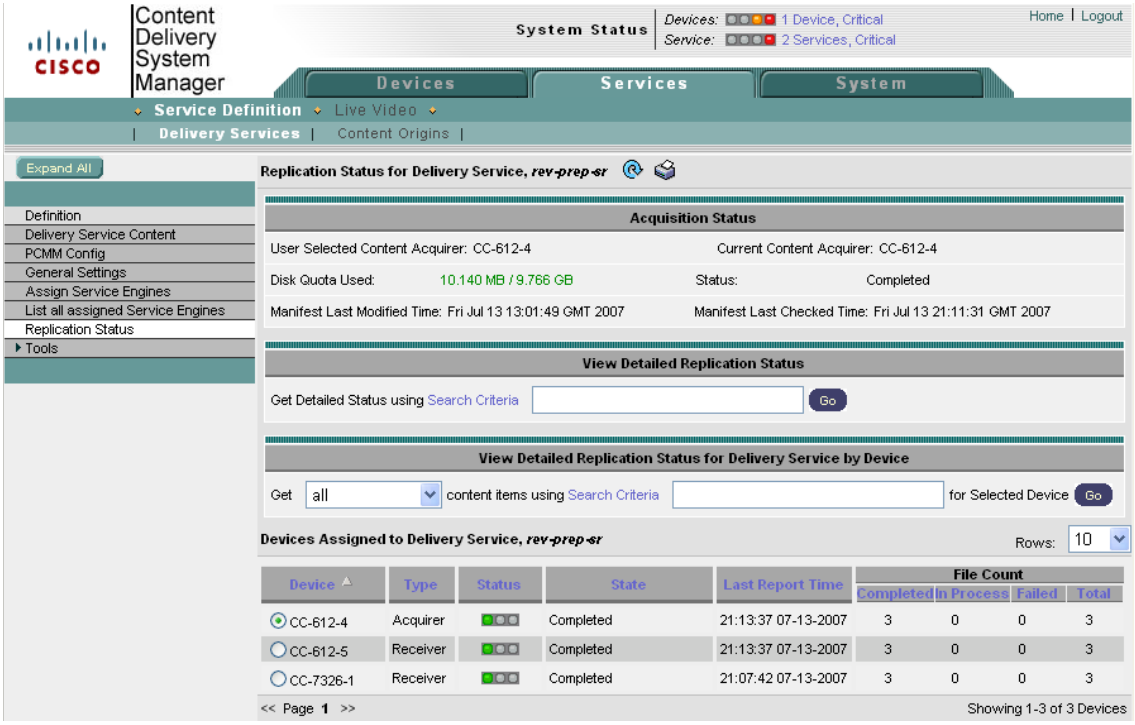

Table 7-16 describes the fields in Acquisition Status section of this page.

This page also allows you to do the following:

- See a detailed view of replication status using search criteria. (See the "Content Replication Status by Delivery Service" section on page 7-23.)
- Query the replication status of content items (by pattern) for a selected SE in the delivery service. (See the "Content Replication Status by Device" section on page 7-25.)

Table 7-16 Replication Status for a Delivery Service

Field	Description
User Selected Content Acquirer	Name of the user-selected Content Acquirer.
Current Content Acquirer	Name of the current Content Acquirer. The current Content Acquirer will be the same as the user-selected Content Acquirer as long as the user-selected one is active; if it fails for any reason, the temporary Content Acquirer becomes the current Content Acquirer.
Disk Quota Used	Amount of available disk space used for the delivery service.
Status	State of the Content Acquirer. (For a description of Content Acquirer states, see Table 7-18.)
Manifest Last Modified Time	Time when the Manifest file was last saved, as recorded on the SE.
Manifest Last Checked Time	Time when the Content Acquirer last checked the Manifest file for changes.

Table 7-17 describes the information about the devices in this delivery service shown at the bottom of the Replication Status page.

Table 7-17 Replication Status for Devices Assigned to a Delivery Service

Field	Description
Device	Name of the SE assigned to the delivery service.
Type	Type of SE: Acquirer, Receiver, or Temporary Acquirer.
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: • Green—No errors encountered. • Yellow—Only minor errors encountered. • Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive SE.
State	State of either the Content Acquirer or receiver SEs. (See Table 7-18 for a description of SE states.)
Last Report Time	Time when the last report from the SE was received by the CDSM. This time stamp uses the CDSM clock.
File Count	
Completed	Number of files that the SE has successfully acquired or received.
In Process	Number of new files to be acquired or replicated. Includes only files for which no acquisition or replication attempts have previously been made.
Failed	For the Content Acquirer: Number of files that failed to be acquired in at least one attempt. For receiver SEs: Number of files that failed to be replicated in at least one attempt. **Note** The failure count for the receiver SE has no relationship to the failure count for the Content Acquirer. If the Content Acquirer fails to replicate an item, the receiver counts this item as "In Process."
Total	Total number of Completed, In Process, and Failed files.

Table 7-18 describes the states of the Content Acquirer or receiver SE.

Table 7-18 Device States

State	Description
Content Acquirer	
Retrieving Manifest	The Content Acquirer is retrieving the Manifest file from the origin server or rechecking the Manifest file for changes.
Processing Manifest	The Content Acquirer has retrieved the Manifest file and is parsing it.
Acquiring Content	The Content Acquirer has processed the Manifest file and is crawling or fetching content.

Table 7-18 Device States (continued)

State	Description
Re-checking Content	The Content Acquirer is checking the content or crawl job freshness.
No Status Reported	No Status Reported can mean: • The Content Acquirer is unreachable for three consecutive polling periods. • The Content Acquirer is offline. • The CDSM has recently restarted and has not yet received a report from the Content Acquirer.
Completed	The Content Acquirer is not in the Retrieving Manifest, Processing Manifest, Acquiring Content, Re-checking Content, or No Status Reported state.
Receiver SE	
Pending Update from Acquirer	The receiver SE is not synchronized with the Content Acquirer.
Replicating	The receiver SE is synchronized with the Content Acquirer and is replicating content.
Completed	The receiver SE has finished replicating all the content with no errors.
Recovering from Failure	The receiver SE has not identified the Content Acquirer. This state occurs during a failover from the Content Acquirer to a temporary Content Acquirer.
No Status Reported	No Status Reported can mean: • The receiver SE is unreachable for three consecutive polling periods. • The receiver SE is offline. • The CDSM has recently restarted and has not yet received a report from the receiver Service Engine.

Content Replication Status by Delivery Service

In the View Detailed Replication Status section of the Replication Status page, enter a search string in the **Get Detailed Status Using** field and click **Go**.

For help on allowed search string characters, click **Search Criteria**.

Use an asterisk (*) to match one or more characters, or a question mark (?) to match only a single character. The criteria are matched against the relative *cdn-url* attribute specified in the <item> tag in the Manifest file. We recommend that you start the search criteria by specifying wildcards such as *.htm or *clip.mpeg.

Figure 7-9 shows the results of a detailed status search for a delivery service.

Figure 7-9 Replication Status for Searched Content Items in a Delivery Service

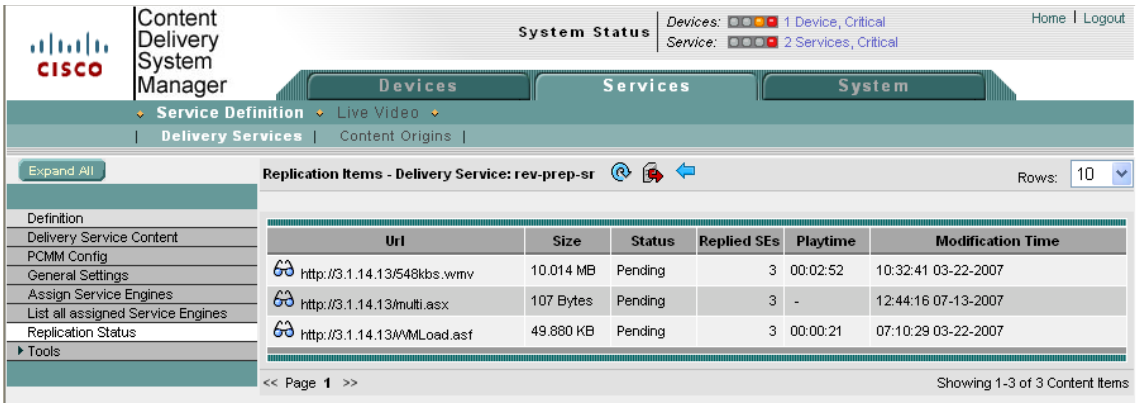

Table 7-19 describes the information displayed for the replication items.

Table 7-19 Replication Status of Items for a Delivery Service

Column Heading	Description
Url	URL of the origin server that stores the content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content in the delivery service. The status is shown as Complete if replication is completed on all SEs assigned to the delivery service.
Replied SEs	Number of SEs that have replicated this item.
Playtime	Duration of playback of the file.
Modification Time	Timestamp of the earliest update for that delivery service from an active SE.

To return to the previous page, click the **Back** icon in the task bar.

To get detailed information about the replication status of the content item, click the **View** icon (eyeglasses) next to the URL. Detailed replication information is displayed (Figure 7-10). This page provides details on the replication status of the content item for every SE in the delivery service. Table 7-20 describes the information on this page.

Figure 7-10 Replication Status for Searched Content Items in a Delivery Service—Detail

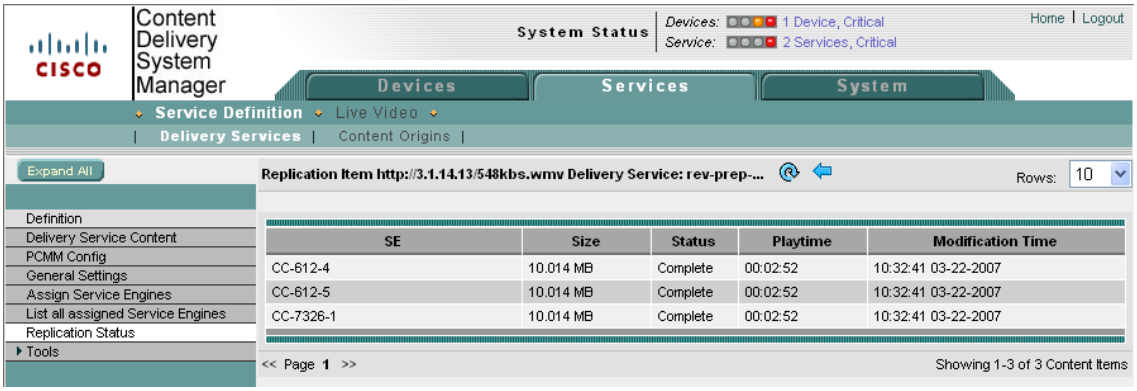

Note The Replication Item page is specifically designed to limit listings to 5000 objects for scalability reasons. These are system limits and not specifically enforced for replication status reporting.

Table 7-20 Replication Status of an Item for All SEs in a Delivery Service

Column Heading	Description
SE	Name of the SE to which the item has been replicated.
Size	Size of the file to be acquired or crawled.
Status	Status of the replication of the content on the SE. Status is shown as Complete if replication is complete on all SEs assigned to the delivery service.
Playtime	Duration of playback of the file.
Modification Time	Timestamp of the latest update for the content item as recorded on the origin server.

To return to the previous page, click the **Back** icon in the task bar.

Content Replication Status by Device

Queries to determine the detailed replication status of a content item trigger extensive CPU cycles and high consumption of memory, because all the SEs assigned to a delivery service need to be polled, and the retrieved replication status is cached in the memory of the CDSM. This results in performance degradation. To optimize the use of memory resources without compromising the need to obtain detailed replication status of a particular content item, you can choose an SE assigned to a delivery service and generate a query.

To view the detailed replication status for a delivery service by device, do the following:

Step 1 From the Replication Status page, in the Devices Assigned to Delivery Service section (see Figure 7-8), click the radio button next to the name of the device that you want to view.

Step 2 In the View Detailed Replication Status for Delivery Service by Device section, do the following:

a. Choose content items (**all**, **replicated**, or **nonreplicated**) from the **Get** drop-down list.

b. In the **Content Items Using** field, enter a string that specifies the type of content items that you want displayed and click **Go**.

Note Use an asterisk (*) to match one or more characters, or a question mark (?) to match only a single character.

The Replication Items page for the selected device is displayed (Figure 7-11). Table 7-21 describes the fields displayed in this page.

Figure 7-11 Replication Items for a Selected Device

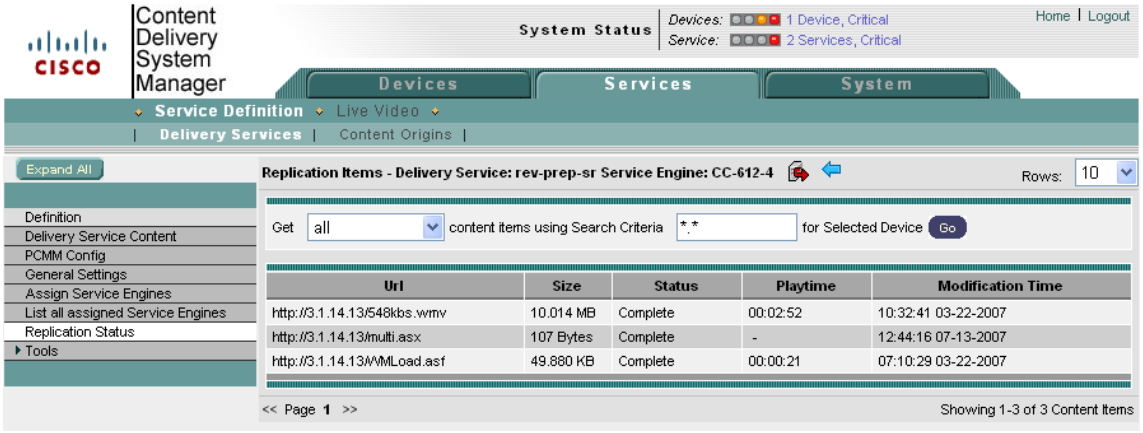

Table 7-21 Replication Status of Items for a Delivery Service by Device

Column Heading	Description
Url	URL of the origin server that stores the content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content for the selected SE.
Playtime	Duration of playback of the file.
Modification Time	Timestamp of the latest update to the content item as recorded on the origin server.

Note When you click the **Force replication information refresh** icon in the task bar, the system displays a dialog box asking you to confirm whether you want to refetch the information from SEs assigned to this delivery service. To continue with the refresh process, click **OK**. You are notified that the request has been queued and are asked to check back later.

Step 3 To refine your search from this window, do the following:

a. Make a choice from the **Get** drop-down list.

b. Enter a search string in the **Content Items Using** field.

c. To retrieve the specified items, click **Go**.

Step 4 To return to the Replication Status page, click the **Back** icon in the task bar.

Viewing Statistics

The Statistics pages track system-wide delivery, replication, and routing traffic in the CDS. You can view statistics on delivery traffic (Movie Streamer, Windows Media, HTTP, and Flash Media) listed by SE or device group. The Routing Statistics page lists client requests and redirects. The Replication Statistics page lists the replication status for all SEs in the CDS, and provides a drill-down to all delivery services for a chosen SE, and all content items associated with that delivery service.

Viewing Service Engines and Device Group Statistics

To view the statistics for all SEs or all device groups, do the following:

Step 1 Choose **Devices > Statistics**. The Statistics page is displayed.

Step 2 Choose **Service Engines** or **Device Groups**, and then choose one of the following content delivery types:

- **Movie Streamer**
- **HTTP**
- **Windows Media**
- Flash Media

Table 7-22 describes each statistic for each content delivery type.

Table 7-22 Service Engine and Device Group Statistics

Statistic	Description
Movie Streamer	
Bandwidth In	Current bandwidth used for input by the Movie Streamer in bits per second.
Bandwidth Out	Current bandwidth used for output by the Movie Streamer in bits per second.
Bytes In	Total bytes received by the Movie Streamer since it was started.
Bytes Out	Total bytes transmitted by the Movie Streamer since it was started.
Packets In	Total packets received by the Movie Streamer since it was started.
Packets Out	Total packets transmitted by the Movie Streamer since it was started.
RTSP Connections	Number of clients currently connected over RTSP.
All Connections	Number of clients connected since startup.
Updated	Timestamp indicating when the statistics were updated.
HTTP	
Requests/Sec	Number of requests per second.
Bytes/Sec	Number of bytes per second.
Request Latency	Average number of seconds per HTTP request.

Table 7-22 Service Engine and Device Group Statistics (continued)

Statistic	Description
Hit Rate	Average number of content items per minute successfully served from the cache of the SE or from all the SEs in the device group.
Updated	Timestamp indicating when the statistics were updated.
Windows Media	
Concurrent Requests	Total number of simultaneous requests the Windows Media Streaming Engine has served.
Kbits/Sec	Total number of kilobits per second served.
Cache Hit Rate	Average number of content items per minute successfully served from the cache of the SE or from all the SEs in the device group
Updated	Timestamp indicating when the statistics were updated.
Flash Media	
Byte/Sec	Total number of bytes per second served.
All Connections	Number of clients connected since startup.
Cache Hit Rate	Average number of content items per minute successfully served from the cache of the SE or from all the SEs in the delivery service.
Updated	Timestamp indicating when the statistics were updated.

Viewing Routing Statistics

To view the routing statistics for SRs, do the following:

Step 1 Choose **Devices > Statistics > Routing Statistics**.

Step 2 Choose one of the following options:

- **Routing Requests**
- **Routing Redirects**

Table 7-23 describes each routing statistic.

Table 7-23 Service Router Statistics

Statistic	Description
Routing Requests	
Total Requests	Total number of content requests received from clients.
HTTP Requests	Number of ASX and traditional HTTP web requests received.
RTSP Requests	Number of RTSP requests received.
RTMP Requests	Number of RTMP requests received.
Updated	Timestamp indicating when the statistics were updated.
Routing Redirects	
Total Requests	Total number of content requests received from clients.
Reqs Redirected	Total number of redirected client requests.
Reqs Not Directed	Total number of client requests not redirected.
Updated	Timestamp indicating when the statistics were updated.

Replication Status

To view system-wide replication status by device, do the following:

Step 1 Choose **Devices > Statistics > Replication Status**. The replication status for all SEs in the CDS is displayed.

Table 7-24 describes the status information displayed on this page.

Table 7-24 Device Replication Status Page

Column Heading	Description
Device	Name of the SE.
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: • Green—No errors encountered. • Yellow—Only minor errors encountered. • Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive SE. (See Table 7-14 for a description of status errors and their corresponding status lights.)
Delivery Service Count	Number of delivery services reporting SEs in a particular state. (See Table 7-18 for a description of SE states.)
Completed	Number of delivery services reporting this SE in a Completed state.

Table 7-24 Device Replication Status Page

Column Heading	Description
In Process	In Process can mean: • Number of delivery services reporting this SE (as a Content Acquirer) in the Retrieving Manifest, Processing Manifest, Acquiring Content, or Re-checking Content state. • Number of delivery services reporting this SE (as a receiver SE) in the Pending Update from Acquirer, Replicating, or Recovering from Failure state.
Failed	Number of delivery services reporting this SE in the Failed or Failed Update state.
Unknown	Number of delivery services reporting this SE in the No Status Reported state.

Step 2 To view the statistics on the delivery services associated with this SE, click **View** (the eyeglasses icon) next to the SE.

The Replication Status for each delivery service that uses the SE to deliver content is displayed. The first column in this table lists the delivery service that uses the SE, the columns that follow list information about the SE's function in the delivery service. For a description of the subsequent columns, see Table 7-17 on page 7-22.

Step 3 To view replication details for the selected delivery service, click the radio button next to the delivery service name.

Step 4 From the **Get** drop-down list, choose the type of items to display (**all**, **replicated**, or **non replicated**).

Step 5 In the **Content Items Using** field, enter a regular expression (such as ***.html**, ***.mpg**, ***.jpg**, or ***.***). Use an asterisk (*) to match one or more characters, and a question mark (?) to match exactly one character.

Step 6 To retrieve the specified items, click **Go**. The Replication Items for Delivery Service page is displayed. Table 7-25 describes the fields displayed in this page.

Note The Replication Items for Delivery Service page is specifically designed to limit listings to 5000 objects for scalability reasons. These are system limits and are not specifically enforced for replication status reporting.

Table 7-25 Replication Status of Items for SEs in a Selected Delivery Service

Column Heading	Description
URL	URL of the origin server that stores the content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content from the Content Acquirer.
Playtime	Duration of playback of the file.
Modification Time	Timestamp of the earliest update for that delivery service from an active SE.

Step 7 To further qualify your search, change the item type from the drop-down list, if you wish, or specify another file type (such as ***.html**, ***.mpg**, or ***.jpg**) in the **Content Items Using** field. To retrieve the specified items, click **Go**.

Step 8 To forcibly refetch the latest content replication information, click the **Force Replication Information Refresh** icon in the task bar. You are asked to confirm whether or not you wish to refetch the information from the SE assigned to the particular delivery service.

Step 9 To continue with the refresh process, click **OK**. You are notified that your request has been sent and prompted to check back after a few minutes.

Step 10 To return to the Replication Status page, click the **Back** button in the task bar.

To view the SE forwarder path for a selected delivery service, click the **View** icon next to the name of the delivery service. To return to the Replication Status page, choose **Replication Status** in the left-panel menu.

Transaction Logs

Transaction logs allow administrators to view the traffic that has passed through the SE. Typical fields in the transaction log are the date and time when a request was made, the URL that was requested, whether it was a cache hit or a cache miss, the type of request, the number of bytes transferred, and the source IP address. For more information about configuring transaction log settings, see the "Configuring Transaction Logs" section on page 4-26.

This section discusses the following topics:

- Transaction Log Formats
- Transaction Logging and NTLM Authentication
- Usage Guidelines for Log Files

Transaction Log Formats

The section discusses the following different logging formats:

- Extended Squid
- Apache
- Custom Format

Extended Squid

The Extended Squid format logs the same fields logged by the Squid-1.1 access log file format.

For details on the Squid-1.1 native log file format, refer to the Squid documentation "Frequently Asked Questions," "Squid Log" section, access.log heading at the following URL:

http://wiki.squid-cache.org/SquidFaq/FaqIndex

The Extended Squid-style log file format is as follows:

Current-Time Time-to-Serve Client-IP Request-Desc/Status-Returned Bytes-Xferred Method URL MIME-Type

An Extended Squid-style log format example looks like this:

```
1012429341.115 100 172.16.100.152 TCP_MISS/302 184 GET http://www.cisco.com/cgi-bin/login
DIRECT/www.cisco.com
```

Apache

The Apache format is the Common Log File (CLF) format defined by the World Wide Web Consortium (W3C) working group. This format is compatible with many industry-standard log tools. For more information, see the W3C Common Log Format website at the following URL:

http://www.w3.org/Daemon/User/Config/Logging.html.

The Apache-style log file format is as follows:

URI SE-IP-address bytes-sent object-size bytes-received method status time-received time-to-serve

An Apache-style log file format example looks like this:

```
http://spcdn-se612-5.se.sanity.spcdn.net/gmedia-0.4gb.wmv SPCDN-SE612-5.spcdn.com
363704065 137 363710748 GET 200 [06/Nov/2007:00:25:32 +0530] 325033158
```

Custom Format

The **transaction-logs format custom** command allows you to use a log format string to log additional fields that are not included in the predefined Extended Squid format or Apache CLF format. The log format string is a string that can contain the tokens listed in Table 7-26 and that mimics the Apache log format string. The log format string can contain literal characters that are copied into the log file. Double backslashes (\\) can be used to represent a literal backslash, and a backslash followed by a single quote (\') can be used to represent a literal single quote. A literal double quote cannot be represented as part of the log format string. The control characters \t and \n can be used to represent a tab and a new line character, respectively.

The following command can be entered to generate the well-known Apache Combined Log Format:

transaction-logs format custom "%t%r %>s %b"

The following transaction log entry example is configured by using the preceding custom format string:

[11/Jan/2003:02:12:44 -0800] "GET http://www.cisco.com/swa/i/site_tour_link.gif HTTP/1.1" 200 3436

Table 7-26 Custom Format Log Format String Values

Format Token	Value
%a	IP address of the requesting client.
%A	IP address of the SE.
%b	Bytes sent, excluding HTTP headers.
%D	Time consumed to serve the request in microseconds
%h	Remote host (IP address of the requesting client is logged).
%H	Request protocol.
%I	Bytes received from the client.
%m	Request method.
%O	Bytes sent to client, including the headers.

Table 7-26 Custom Format Log Format String Values (continued)

Format Token	Value
%q	Query string (which is preceded by a question mark (?) if a query string exists; otherwise, it is an empty string).
%r	First line of the request.
%>s	Status. The translog code always returns the HTTP response code for the request.
%t	Time in common log time format (or standard English format).
%T	Time consumed to serve the request in seconds (a floating point number with 3 decimal places).
%U	URL path requested, not including query strings.
%V	Value of the host request header field reported if the host appeared in the request. If the host did not appear in the host request header, the IP address of the server specified in the URL is reported.
%X	Connection status when the response is completed.

Transaction Logging and NTLM Authentication

If your device is configured for NT LAN Manager (NTLM) authentication and uses the Apache-style or Extended Squid-style format, you can record the Windows domain name and username in the "authenticated username" field of the transaction log. If the domain name is available, both the domain name and the username are recorded in the "authenticated username" field, in the form domain\username. If only the username is available, only the username is recorded in the "authenticated username" field. If neither a domain name nor a username is available, a "-" (hyphen) is recorded in the field.

Usage Guidelines for Log Files

This section provides some guidelines for working with log files, and includes the following topics:

- Working Logs
- Archive Working Log
- Exporting Log Files
- Windows Media Transaction Logging

Working Logs

Transaction logs are located in the /local/loca1/logs directory. Each component has one or more directories, depending on its configuration.

There is a working log file in each directory, which is a symbol link, linking to the current working log file.

The log files are logged to a working log on the local disk as follows:

- WMT logs are logged to a working log on the local disk in /local1/logs/export/working.log
- Movie Streamer logs are logged to a working log on the local disk in /local1/logs/movie-streamer/working.log

- Flash Media Streaming logs are logged to a working log on the local disk in /local1/logs/fms_access/working.log and /local1/logs/fms_authorization/working.log

Note For Movie Streamer, client requests that join the multicast group do not appear in the transaction log because multicast clients do not contact the server.

Archive Working Log

You can specify the interval at which the working log should be cleared by moving the data to an archive log. The archive log files are located on the local disk in the /local1/logs/ directory.

Archive files can be configured by time interval and file size. If one of the criteria is met, a log rotation occurs. You can specify the maximum number of old logs kept on disk.

Because multiple archive files are saved, the filename includes the timestamp when the file was archived. Because the files can be exported to an FTP/SFTP server, the filename also contains the IP address of the SE.

The archive filenames use this format: modulename_IPADDRESS_YYYYMMDD_HHMMSS.

For example, fms_access_10.74.61.130_20070913_080051 is the filename for the archive of the fms_access log.

Note The IP address used in the archived filename is not necessarily the primary interface of the SE, the transaction log function decides on which IP address to use in creating the archive name.

Exporting Log Files

To facilitate the post-processing of cache log files, you can export transaction logs to an external host. This feature allows log files to be automatically exported by FTP to an external host at configurable intervals. The username and password used for FTP are configurable, as is the directory to which the log files are uploaded.

The log files automatically have a filename that uses the *<type>_<ipaddr>_yyyymmdd_hhmmss* format, where:

- *<type>* represents the type of log file, with *selog* for cache logs such as HTTP, HTTPS, and FTP, and *mms_export* for Windows Media Technologies (WMT) logs.
- *<ipaddr>* represents the SE IP address.
- *yyyymmdd_hhmmss* represents the date and time when the log was archived for export.

Exporting Transaction Logs to External FTP Servers

To export transaction logs to an FTP server, you must first enable exporting of transaction logs and then configure the FTP or secure FTP (SFTP) server parameters. This feature can support up to four FTP servers. The following information is required for each target FTP server:

- Server IP address or the hostname

  The SE translates the hostname with a DNS lookup and then stores the IP address in the configuration.

- FTP user login and user password

- Path of the directory where transferred files are written

  Use a fully qualified path or a relative path for the user login. The user must have write permission to the directory.

You can also compress archived log files into gzip format before exporting them to external FTP servers. The compressed filename has a .gz extension. This compression feature uses less disk space than that required for noncompressed archived files on both the SE and the FTP export server and also requires less bandwidth during export because of the smaller size of the files to be exported.

Restarting Export After Receiving a Permanent Error from the External FTP Server

When an FTP server returns a permanent error to the SE, the archive transaction logs are no longer exported to that server. You must re-enter the SE transaction log export parameters for the misconfigured server to clear the error condition.

A permanent error (Permanent Negative Completion Reply, RFC 959) occurs when the FTP command to the server cannot be accepted, and the action does not take place. Permanent errors can be caused by invalid user logins, invalid user passwords, and attempts to access directories with insufficient permissions or directories that do not exist.

Exporting Transaction Logs to External SFTP Servers

You can also export transaction logs to a Secure File Transfer Protocol (SFTP) server. You must first enable the feature and configure the SFTP server parameters. The following information is required for each target SFTP server:

- SFTP server IP address or the hostname

  The SE translates the hostname with a DNS lookup and then stores the IP address in the configuration.

- SFTP user login and user password
- Path of the directory where transferred files are written

  Use a fully qualified path or a relative path for the user login. The user must have write permission to the directory.

To enable this feature, enter the **sshd allow-non-admin-users** command on the SE. If this feature is enabled, the output of the **show running-config** EXEC command shows that this feature is enabled on the SE.

Windows Media Transaction Logging

The following logging formats are supported for Windows Media transaction logging:

- Standard Windows Media Services Version 4.1
- Extended Windows Media Services Version 4.1
- Standard Windows Media Services Version 9.0
- Extended Windows Media Services Version 9.0

The extended versions of the logging formats contain additional fields that are SE specific (For example, the SE-action field specifies a cache hit or miss, and the SE-bytes field specifies the number of bytes that were sent from the SE.)

The SE's transaction logging format for Windows Media streaming is consistent with that of the Windows Media Services and the World Wide Web Consortium (W3C)-compliant log format. A log line is written for every stream accessed by the client. The location of the log is not configurable. These logs can be exported using FTP. When transaction logging is enabled, daemons create a separate *working.log* file in /local1/logs/export for WMT transactions.

All client information in the transaction logs is sent to the origin server by default.

Note All WMT playable contents can be delivered by either HTTP or RTSP, based on the request. Any content that is cached by the WMT is stored using the RTSP scheme, regardless of whether the content was cached due to an HTTP or RTSP request. Therefore, in the **show** command, the content displays as RTSP.

Log Formats Accepted by Windows Media Services 9

Windows Media Players connect to a Windows Media server using the following protocols:

- Windows Media Players earlier than Version 9.0 use HTTP/1.0 or the MMS protocol.
- Windows Media Player Version 9.0 uses HTTP/1.1 and RTSP.

Depending on the version of the Windows Media Player, logs are sent in different formats, such as text, binary, or Extensible Markup Language (XML). Table 7-27 describes the log formats accepted by Windows Media Services Version 9.0.

Table 7-27 Windows Media Services Version 9.0 Log Formats

Protocol	Player and Distributor	Log Type
HTTP/1.0	Windows Media Player earlier than Version 9.0 SE (caching and proxy server) is running Windows Media Services Version 9.0 and streaming from a Windows Media server that is running Windows Media Services Version 4.1	World Wide Web Consortium (W3C) standard space-delimited text log
MMS	Windows Media Player earlier than Version 9.0	Binary structure log
HTTP/1.1	Windows Media Player Version 9.0 Distribution server is running Windows Media Services Version 9.0 SE (caching and proxy server) is running Windows Media Services Version 9.0	XML structure log
RTSP	Windows Media Player Version 9.0 Distribution server is running Windows Media Services Version 9.0 SE (caching and proxy server) is running Windows Media Services Version 9.0	XML structure log

Note Extensible Markup Language (XML) logging for MMS-over-HTTP and MMS-over-RTSP (RTSP over Windows Media Services Version 9.0) is supported. The posted XML log file from the Windows Media Player to the SE (Windows Media server) can be parsed and saved to the normal Windows Media transaction logs that are stored on the SE.

CHAPTER 8

Maintaining the Internet Streamer CDS

This chapter explains how to perform common administrative tasks including updating system software, hard disk drive maintenance, and rebooting and deleting devices. The following major topics are covered:

- Software Upgrade, page 8-1
- Rebooting Devices, page 8-9
- Deleting a Device, page 8-10
- Replacing a Device, page 8-12
- Backup and Recovery Procedures, page 8-15
- Disk Maintenance, page 8-26

For information about database maintenance, see the "Scheduling Database Maintenance" section on page 4-55.

Software Upgrade

The software upgrade section covers the following topics:

- Getting a Software File from Cisco.com
- Finding the Software Version of the Devices
- Configuring the Software Image Settings
- Upgrading the Software
- Software Upgrades by Device

Getting a Software File from Cisco.com

To get a software file from Cisco.com, do the following:

Step 1 Launch your web browser and enter the following URL:

http://www.cisco.com/kobayashi/sw-center/sw-video.shtml

The Log In page is displayed.

Step 2 Log in to Cisco.com using your designated username and password. The Video and Content Delivery page is displayed, listing the available software products.

Step 3 Click **Cisco Content Delivery Systems (CDS)**. The Downloads page is displayed.

Step 4 Click the **Cisco Content Delivery Applications** folder to expand it, and click the **Cisco Internet Streamer Application**. The page refreshes and the software releases are displayed.

Step 5 Click the software release you want. The page refreshes and the software image files are displayed.

Step 6 Click the link for the software image file you want.

- If this is the first time you have downloaded a file from Cisco.com, the Cisco Systems Inc., Encryption Software Usage Handling and Distribution Policy is displayed. Read the policy, fill in the unfilled fields, and click **Accept**.
- If you previously filled out the Encryption Software Usage and Handling and Distribution form, the form does not display again.

The Download page is displayed with the information about the software image file and a Download link.

Step 7 Click **Download**. The Cisco End User Software License Agreement is displayed.

Step 8 Read the agreement and click **Agree**. The File Download dialog box is displayed.

Step 9 Click **Save**. The Save As dialog box is displayed.

Step 10 Navigate to the location where you want to save the file and click **Save**. The file downloads.

Pre-positioning a Software File

A software file is pre-positioned in the same manner as any other content item. Pre-positioning allows you to conserve bandwidth usage across the WAN and avoid congesting your network during updates. The software file is fetched one time from the origin server, replicated across your network, and stored in Service Engine caches in your LAN.

To pre-position a software file, you must complete the following tasks:

- Define a delivery service.
- Assign devices to the delivery service.
- Define the software file that you want to pre-position by using a Manifest file or the CDSM delivery service content page.
- Check the device replication status.

See Chapter 5, “Configuring Services” for more information.

Note Only Service Engines that are assigned to the delivery service can be updated using pre-positioned software files. Service Routers and CDSMs do not have pre-positioned content; therefore, you cannot use the pre-positioned method for device updates for these devices.

Sample Manifest File to Pre-position a Software File

You can use the following sample Manifest file to pre-position a software file by replacing the URL with a valid software file URL:

```
<CdnManifest>
<item src="http://your-web-server.com/folder/upgrade.bin" />
</CdnManifest>
```

The server name or IP address of the URL in the Manifest file (and in the Software File URL field in the Software File Settings page must match either the Origin Server field or the Service Router Domain Name field in the Content Origin page).

Finding the Software Version of the Devices

The CDSM Home page gives a brief summary of the software versions in use on all the devices in the CDS network.

To view the software version running on a particular device, choose **Devices > Devices**. The Devices Table page displays the software version for each device listed.

Clicking the **Edit** icon next to the device name in the Devices Table page displays the Device home page, which shows the software version for that device.

Note The software version is not upgraded until a software upgrade has been successfully completed. If a software upgrade is in progress, the version number displayed is the base version, not the upgraded version number.

Configuring the Software Image Settings

To upgrade your software from CDS 2.0.x or a later release, you must first configure the software image settings.

To configure the software image settings, do the following:

Step 1 Choose **System > Software Image Management**. The Software Files Table page is displayed.

Step 2 Click the **Create New** icon in the task bar. The Software Image page is displayed (Figure 8-1).

Figure 8-1 Software Image Page

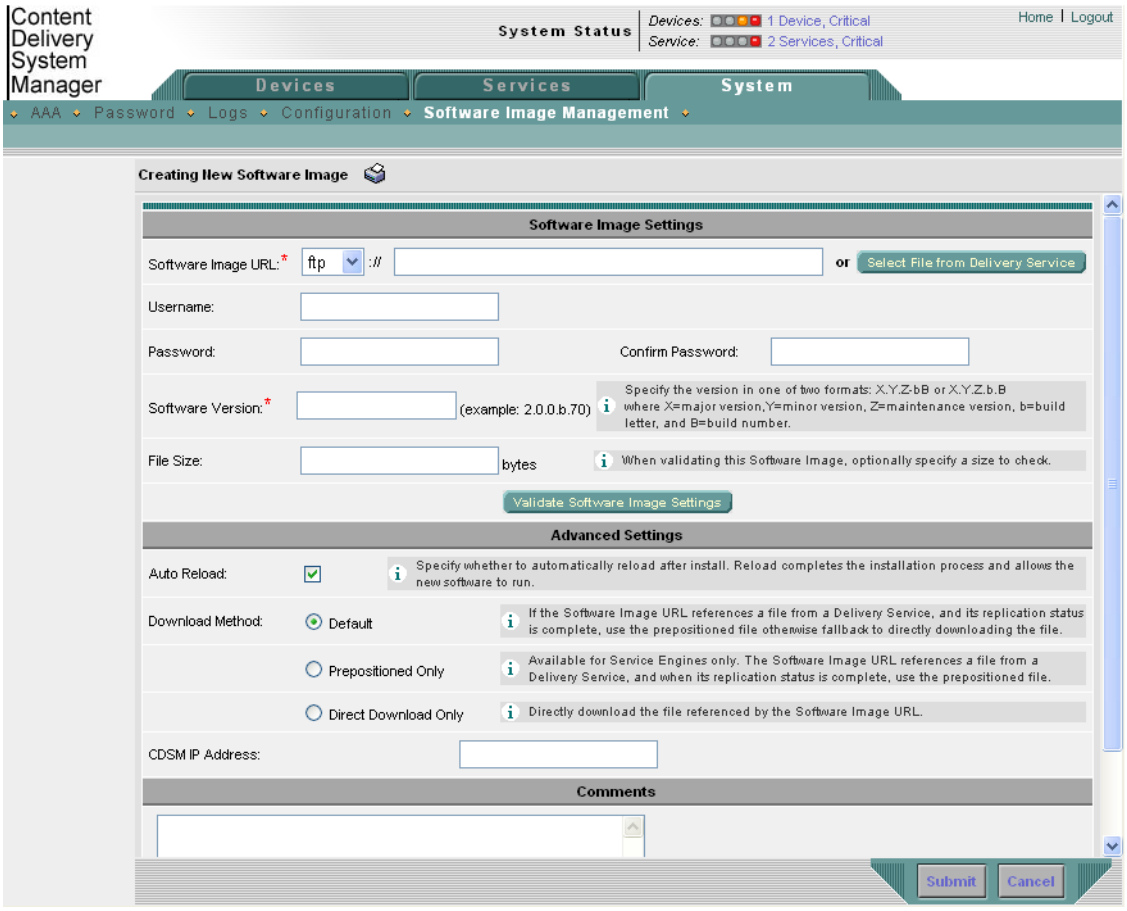
Content Delivery System Manager
System Status
Devices: 1 Device, Critical
Service: 2 Services, Critical
Home | Logout
Devices
Services
System
AAA
Password
Logs
Configuration
Software Image Management
Creating New Software Image
Software Image Settings
Software Image URL:*
ftp
://
or
Select File from Delivery Service
Username:
Password:
Confirm Password:
Software Version:*
(example: 2.0.0.b.70)
Specify the version in one of two formats: X.Y.Z-bB or X.Y.Z.b.B where X=major version,Y=minor version, Z=maintenance version, b=build letter, and B=build number.
File Size:
bytes
When validating this Software Image, optionally specify a size to check.
Validate Software Image Settings
Advanced Settings
Auto Reload:
Specify whether to automatically reload after install. Reload completes the installation process and allows the new software to run.
Download Method:
Default
If the Software Image URL references a file from a Delivery Service, and its replication status is complete, use the prepositioned file otherwise fallback to directly downloading the file.
Prepositioned Only
Available for Service Engines only. The Software Image URL references a file from a Delivery Service, and when its replication status is complete, use the prepositioned file.
Direct Download Only
Directly download the file referenced by the Software Image URL.
CDSM IP Address:
Comments
Submit
Cancel
211837

Step 3 In the **Software Image URL** field, enter the URL for the .bin software file that you downloaded from Cisco.com.

a. Choose a protocol (**http** or **ftp**) from the drop-down list.

b. Enter the URL of the software file; for example, a valid URL might look like this:

http://internal.mysite.com/cds/*CDS-2.x.x-K9*.bin

In this URL, *CDS-2.x.x-K9* is the name of the software upgrade file. (The filename might include the version number.)

Note If you are using a pre-positioned software file and you are entering the URL manually (rather than using the **Select File from Delivery Service** option), the server name or IP address of the URL in the Software Image URL field must match either the Origin Server field or the Service Routing Domain Name field in the Content Origin page of the delivery service. This is not a requirement if you are downloading the software file directly from the origin server. (See the "Pre-positioning a Software File" section on page 8-2 for details.)

Alternatively, click **Select File from Delivery Service**. A separate window displays that allows you to choose a delivery service, set criteria, search the delivery service, and select the software file that you want to use for the software upgrade. (You must first pre-position the software file in the delivery service. See the "Pre-positioning a Software File" section on page 8-2.)

Step 4 If your server requires user login authentication, enter your username in the **Username** field and enter your login password in the **Password** field. Enter the same password in the **Confirm Password** field.

Step 5 Enter the software version number in the **Software Version** field. You can copy this number from the version portion of the software filename in the software file URL.

Specify the version in one of two formats: X.Y.Z-bB or X.Y.Z.b.B, where X = major version, Y = minor version, Z = maintenance version, b = build letter, and B = build number.

Step 6 If you want the size of the software file considered during validation, enter a file size (in bytes) in the **File Size** field. If you leave this field blank, the URL is checked without regard to the software file size.

Step 7 To validate the Software Image URL, Username, and Password fields, click the **Validate Software Image** button.

When you click the **Validate Software Image Settings** button, the following occurs:

- The software file URL is resolved.
- A connection to the software file URL is established using the username and password, if specified.
- If a file size is specified, the actual size of the software file is obtained and compared against the value in the File Size field.
- A message is returned, indicating success or errors encountered.

Step 8 In the Advanced Settings section, check the **Auto Reload** check box to automatically reload a device when you upgrade the software.

Step 9 If you want, you can choose one of three download methods:

- **Default**—Uses pre-positioned content but always falls back to direct download.
- **Prepositioned Only**—Uses the local file copy if the software file URL references pre-positioned content and its replication status is complete.
- **Direct Download Only**—Directly downloads the file using the software file URL.

Note If you choose **Prepositioned Only**, the software file settings that you define in this page cannot be used to upgrade a CDSM or an SR, because these devices do not have pre-positioned content.

Step 10 For downgrades only, specify the CDSM IP address to be used for device registration in the **CDSM IP Address** field.

The CDSM IP Address field is the IP address of a CDSM after the software is downgraded. (This field is optional and only applies for downgrades.) After the downgrade, the SE registers with the CDSM with the IP address specified in this field.

Step 11 Click **Submit**.

To delete a software file, click the **Trash** icon in the task bar.

Caution If your browser is configured to save the username and password for the CDSM, the browser will auto-populate the Username and Password fields in the Software Image page. You must clear these fields before you click **Submit**.

The software file that you want to use is now registered with the CDSM. When you perform the software upgrade or downgrade, the URL that you just registered becomes one of the choices available in the Update Software page. (See the "Upgrading the Software" section on page 8-6.)

Upgrading the Software

When upgrading software in your CDS network, begin with Service Engines and Service Routers before upgrading the CDSM. The CDSM reboots at the conclusion of the upgrade procedure, causing you to temporarily lose contact with the device and the user interface. After the CDSM has upgraded its software and rebooted, it may be unable to communicate with devices running different versions of the CDS software.

Caution Primary and standby CDSMs must be running the same version of CDS software. If they are not, the standby CDSM detects this and will not process any configuration updates it receives from the primary CDSM. You need to upgrade your standby CDSM first, and then upgrade your primary CDSM. We also recommend that you create a database backup for the primary CDSM and copy the database backup file to a safe place before you upgrade the software.

Note For software downgrades, you need to do the following:

1. If you are using the CDSM GUI, downgrade the standby CDSM first, followed by the primary CDSM.

   If you are using the CLI, downgrade the primary CDSM first , followed by the standby CDSM.

2. After downgrading the primary and standby CDSMs, using the CLI, log in to each CDSM and run the following commands:

   ```
   cms database downgrade
   cms enable
   ```

3. Downgrade the software on the Service Routers, followed by the Service Engines.

Caution To upgrade the software image on a server, you first need to offload a server for maintenance. Once the server has been fully off loaded, you can upgrade the software. After updating the software, uncheck the **Server Offload** check box to allow the server to receive client requests from the Service Router. See the Server Offload field in Table 4-2 on page 4-8 for more information.

Interoperability Considerations

In general, a CDS network will be upgraded gradually, so that your network might consist of nodes with different software versions for some time. You can expect the following behavior during an upgrade or downgrade of your network:

- The CDS network continues to operate with mixed versions up to one major or minor version difference in a deployed solution.
- New features that depend on device cooperation might not be fully functional until the CDS network upgrade is complete, but no existing features will be affected.

- While being upgraded, a node will be unavailable for a short time.
- All nodes, other than the node being upgraded, continue to operate at full capacity. The availability of other nodes is not affected during an upgrade.
- Content is preserved during an upgrade or downgrade unless you remove a delivery service.
- All logs are preserved during an upgrade or downgrade, unless you change the disk configuration. Anytime disk space is reconfigured, the logs are automatically removed.

We strongly recommend that you upgrade your CDS network devices in the following order:

1. Multicast sender Service Engines
2. Multicast receiver Service Engines
3. Non-Content Acquirer Service Engines
4. Content Acquirer Service Engines
5. Service Routers
6. Standby CDSMs (Upgrade before primary when using the GUI only.)
7. Primary CDSM

Note When you upgrade CDSMs using the CLI, we recommend that you upgrade your primary CDSM first, and then upgrade your standby CDSM. Primary and standby CDSMs must be operating with exactly the same software release as each other for failover to be successful.

Upgrading Software by Device Groups

Note This procedure is for Service Engines only. Service Routers and CDSMs cannot be associated with device groups.

To upgrade your software on multiple Service Engines, do the following:

Step 1 Choose **Devices > Device Groups**. The Device Groups Table page is displayed.

Step 2 Click the **Edit** icon next to the name of the device group that you want to upgrade. The Device Group page is displayed.

Step 3 From the left-panel menu, choose **Software Update**. The Software Update for Device Group page is displayed.

Step 4 Choose the software file URL from the Software File URL list by clicking the radio button next to the filename.

Step 5 Click **Submit**.

To view progress on an upgrade, go to the Devices Table page (**Devices > Devices**). Software upgrade status messages are displayed in the Software Version column. These intermediate messages are also written to the system log on the Service Engines. See Table 8-1 for a description of upgrade status messages.

Table 8-1 Upgrade Status Messages

Upgrade Status Message	Condition
Pending	The request has yet to be sent from the CDSM to the device, or receipt of the request has yet to be acknowledged by the device.
Downloading	The download method for the software file is being determined.
Proceeding with Pre-positioned Download	The download method for the software file is detected as pre-positioned. Proceeding with download of a pre-positioned software file.
Proceeding with Download	The download method for the software file is detected as direct download. Proceeding with the request for direct download of the software file.
Download in Progress (Completed …)	Direct download of the software file is being processed. "Completed" indicates the number of megabytes processed.
Download Successful	The direct download of the software file has been successful.
Download Failed	The direct download of the software file cannot be processed. Further troubleshooting is required; see the device system message log.
Proceeding with Flash Write	A request has been made to write the software file to the device flash memory.
Flash Write in Progress (Completed …)	The write of the device flash memory is being processed. "Completed" indicates the number of megabytes processed.
Flash Write Successful	The flash write of the software file has been successful.
Reloading	A request to reload the device has been made in order to complete the software upgrade. The device may be offline for several minutes.
Reload Needed	A request to reload the device has not been made. The device must be reloaded manually to complete the software upgrade.
Canceled	The software upgrade request was interrupted, or a previous software upgrade request was bypassed from the CLI.
Update Failed	The software upgrade could not be completed. Troubleshooting is required; see the device system message log.

Software Upgrades by Device

Use this upgrade procedure for Service Routers and CDSMs. You can also use this upgrade procedure to upgrade Service Engines one at a time.

To upgrade your software on a single device, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed.

Step 2 Click the **Edit** icon of the device that you want to upgrade. The Device home page is displayed.

Step 3 Verify that the device is not already running the version that you plan to upgrade to, and that the current version has an upgrade path to the version that you plan to upgrade to.

Step 4 Click **Update Software**. The Software Update page is displayed.

Step 5 Choose the software file URL from the Software Files list by clicking the radio button next to the filename.

Step 6 Click **Submit**, and then click **OK** to confirm your decision.

The Devices Table page is displayed again. You can monitor the progress of your upgrade from this page.

Software upgrade status messages are displayed in the Software Version column. These intermediate messages are also written to the system log on the Service Engines. See Table 8-1 for a description of upgrade status messages.

Rebooting Devices

You can reboot a device or device group. The CDSM performs a controlled shutdown of all devices and then restarts the operating system on each device.

To reboot an individual device, do the following:

Step 1 Choose **Devices > Devices**.

Step 2 Click the **Edit** icon next to the device name that you want to reboot. The Device home page is displayed

Step 3 In the task bar, click the **Reload** icon. You are prompted to confirm your decision.

Step 4 To begin rebooting the device, click **OK**.

To reboot an entire device group, follow these steps:

Step 1 Choose **Devices > Device Groups**.

Step 2 Click the **Edit** icon next to the name of the device group that you want to reboot. The Device Group page is displayed.

Step 3 In the task bar, click the **Reboot All Devices in Device Group** icon. You are prompted to confirm your decision.

Step 4 To begin rebooting each SE in the device group, click **OK**.

Deleting a Device

You can delete a device if the device is experiencing unresolvable problems or when its network address or configuration has changed and you need to add the device back to the CDS network using its new address and configuration information.

Caution If you delete the only SR in your CDS network, you are removing the ability of your CDS network to fill user requests.
When you delete an SE from the CDS network, you are removing that device and the content it contains from the routing scheme that the CDS uses to fill user requests. Although the CDS routes requests around SEs that are busy, offline, or missing, removing an SE may affect the speed at which the CDS network can serve user requests.

Note You cannot delete an SE if it is the only device assigned to a location that is designated as the root location (Content Acquirer) for a delivery service and there are other SEs associated with the delivery service. You can delete the Content Acquirer for a delivery service if the Content Acquirer is the only SE associated with that delivery service. However, deleting the only SE in a delivery service makes the delivery service unable to deliver content. If you receive an error message referencing the Content Acquirer for a delivery service, add more SEs to that location, or change the root location by choosing an SE in a different location to be the Content Acquirer for the delivery service before attempting to delete the SE again.

Removing the device from the CDS network involves using the CLI to shut down CDS network services and deregister the node. If you are removing the device because of hardware failure and it cannot be accessed through its CLI, you can remove the device by using the CDSM; however, the device continues to store its registration information until you deregister it by using the CLI.

Before a device can be removed from the CDS network, the following conditions must be met:

- The device must have been activated in the CDSM.
- The CDSM must be operating.
- The device must have the correct CDSM IP address or hostname configured.
- The CDSM IP address or hostname must be that of the primary CDSM.

Deleting a device from the CDS network involves using the CLI to remove the registration information from the device itself and removing the registration record from the CDSM.

Note Do not use the CDSM to delete a device while the device is still active and registered. The CDSM delete feature removes only the device's registration record from the CDSM; it does not deregister the device. The device retains its registration information and continues to contact the CDSM; however, the CDSM no longer recognizes the device.

If for some reason the CDSM loses the registration record of a device, use the **cms deregister force** command on the device to remove all of its registration information. Then use the **cms enable** command to reregister the device with the CDSM as though it were a new node in the CDS network.

To remove and deregister a device, do the following:

Step 1 Open an SSH session to the device CLI.

Step 2 In global configuration mode, enter the **no cms enable** command.

```
SE# configure
SE(config)# no cms enable
```

Note Issuing the **no cms enable** command does not disable acquisition and distribution services on the device; however, issuing the **cms deregister** command does. The **cms deregister** command disables the CMS, all acquisition and distribution services, and all routing communications to and from this device.

Step 3 In EXEC mode, enter the **cms deregister** command.

```
SE(config)# exit
SE# cms deregister
```

Note The **cms deregister** command cleans up the database automatically. You do not need to use the **cms database delete** command.

Step 4 If for some reason the deregistration fails, you can force the deregistration by using the **cms deregister force** command.

```
SE# cms deregister force
```

Note Ignore any messages stating that the deregistration failed. The **cms deregister force** command forces the deregistration to continue.

Step 5 To add the device back into the CDS network, reregister the device with the CDSM by using the **cms enable** command in global configuration mode.

```
SE# configure
SE(config)# cms enable
```

In case of a hardware failure, you might need to remove the device from the CDS network routing scheme by using the CDSM.

Before a device can be removed from the CDS network through the CDSM, the following conditions must be met:

- The device must have been activated in the CDSM.
- The CDSM must be running.
- The device must have the correct CDSM IP address or hostname configured.
- The CDSM IP address or hostname must point to the primary CDSM.

To delete a device using the CDSM, do the following:

Step 1 Choose **Devices > Devices**. The Devices Table page is displayed. The online status of the device is listed in the Status column.

Step 2 Click the **Edit** icon next to the device name you want to delete. The Device home page is displayed.

Step 3 In the task bar, click the **Delete Device** (**Trash**) icon. You are prompted to confirm your decision.

Step 4 To execute your request, click **OK**. The device is removed from the CDSM.

Step 5 If possible, access the device CLI to deregister the device.

Step 6 In the CLI, enter the **cms deregister force** command.

Note You must use the **cms deregister force** command after deleting a device in the CDSM. This is because once the device has been deleted, the CDSM no longer has a record of the device.

Step 7 To add the device back in to the CDS network, reregister the device with the CDSM by using the **cms enable** command in global configuration mode.

Deleting a Warm Standby CDSM

You can delete a warm standby CDSM from the CDS network at any point after you have registered the device and before the device has come online as the primary CDSM. Once the device has been called into use as the primary CDSM, however, you cannot delete it by using the CDSM.

Delete a warm standby CDSM when the device is experiencing unresolvable problems or when its network address or configuration has changed and you need to add the device back to the CDS network by using its new address and configuration information.

To delete a warm standby CDSM, do the following:

Step 1 Log in directly to the CDSM CLI, and enter the **cms deregister** command.

If for some reason the deregistration fails, you can force the deregistration by using the **cms deregister force** command.

Step 2 From the CDSM GUI, choose **Devices > Devices**.

The browser refreshes, listing the CDSMs on your CDS network. The warm standby CDSM is identified as *Standby*.

Step 3 Click the **Edit** icon next to the name of the warm standby CDSM. The Device home page is displayed.

Step 4 From the left-pane menu, choose **Device Activation**. The Activation page is displayed.

Step 5 In the task bar, click the **Trash** icon. You are prompted to confirm your decision.

Step 6 To execute your request, click **OK**.

Replacing a Device

The procedure to replace a device in the CDS is different depending on the type of the device being replaced. This section covers the following procedures:

- Replacing a CDSM
- Replacing an SE or SR

Replacing a CDSM

To replace a CDSM in a CDS you must first add the new CDSM into the network as a standby CDSM. For procedural information, see the "Configuring Primary and Standby CDSMs" section on page 3-8.

Note The primary and standby CDSMs must be running the same version of software. You must first add the new CDSM with the same version as the existing CDSM. Once the standby CDSM has been added, you must wait at least two polling intervals (10 minutes) for the databases to synchronize before you can begin the upgrade procedure.

Note After you have activated the standby CDSM using the primary CDSM web interface and the device shows as online in the Devices Table page, wait at least two polling intervals (10 minutes) before changing roles to ensure that the standby CDSM has a record of the most recent configuration changes.

To promote the standby CDSM to primary, first stop the primary CDSM using the **cdsm role standby** command. For procedural information, see the "Changing a Standby to a Primary CDSM" section on page 3-9.

After the primary CDSM has been stopped, and the standby CDSM has taken the role of primary, wait at least two polling intervals (10 minutes) before logging in to the new primary CDSM. The new primary CDSM is accessible by entering the IP address of the CDSM with port 8443 in a web browser. For example, if the IP address of your CDSM is 192.168.0.236, enter **https://192.168.0.236:8443**.

It is now safe to deactivate the old primary CDSM in the CDSM web interface and remove it from the CDS network.

Note Do not try to take a back up of the old CDSM database and restore it on the new CDSM. This may lead to problematic issues.

Replacing an SE or SR

Note If you replace a Content Acquirer with an SE that was not previously assigned to the delivery service, all content is reacquired and old content is deleted.

Note To prevent the reacquisition of content when replacing a Content Acquirer, make one of the receiver SEs in the same delivery service the replacement Content Acquirer. Add the new SE as a receiver SE, wait until replication is complete for the newly added SE, and then designate it as the Content Acquirer. When you replace a Content Acquirer in this manner, the SEs in the delivery service synchronize with the new Content Acquirer through the metadata poll. Content is not redistributed to the other SEs in the delivery service unless the content has changed since the last metadata poll.

To replace an SE or SR, do the following:

Step 1 Open an SSH session to the device being replaced.

Step 2 In global configuration mode, enter the **no cms enable** command to disable CMS on the device that needs to be replaced.

```
SE# configure
SE(config)# no cms enable
```

Step 3 From the CDSM, choose **Devices > Devices > Device Activation**. The Device Activation page is displayed.

Step 4 Uncheck the **Activate** check box and click **Submit**. The page refreshes and displays a **Replaceable** check box.

Step 5 Check the **Replaceable** check box and click **Submit**.

Step 6 Choose **System > Configuration > System Properties**. The System Properties page is displayed.

Step 7 Click the edit icon next to the **System.devivce.recovery.key** property. The Modify Config Property page is displayed.

Step 8 In the **Value** field, enter a key and click **Submit**. The default value is default.

Step 9 Follow the instructions for configuring a device using the setup utility. The instructions can be found in the *Cisco Content Delivery Engine 100/200/300/400 Hardware Installation Guide*.

Note The replacement device must be the same hardware model as that of the device being replaced.

a. When prompted by the setup utility, configure the basic network settings.

b. When prompted by the setup utility for the hostname of the new device, use the same hostname of the device being replaced. For example, if the old device has a hostname of “SE1,” the new device must have a hostname of “SE1.”

c. When prompted by the setup utility for the IP address of the CDSM, enter the IP address of the CDSM.

Step 10 Open an SSH session to the new device.

Step 11 In EXEC mode, enter the **cms recover identity** command with the key parameter you set in Step 8.

```
SE# cms recover identity <key>
```

On successful registration to the CDSM, a message similar to the following is displayed:

```
DT-7326-4#cms recover identity sr
Registering this node as Service Router...
Sending identity recovery request with key sr
Node successfully registered with id CrConfig_291
Registration complete.
```

Step 12 Register the device with the CDSM by using the **cms enable** command in global configuration mode.

```
SE# configure
SE(config)# cms enable
```

Step 13 From the CDSM, choose **Devices > Devices > Device Activation**. The Device Activation page is displayed.

Step 14 Check the **Activate** check box and click **Submit**.

After a few minutes, approximately two polling intervals, the device status shows online and all configurations (delivery service assignments, programs, and so on) are the same as those on the device that was replaced.

Step 15 Once the new device is up and running, as noted by the online status, the old device can be removed from the CDS network.

Backup and Recovery Procedures

This section provides CDSM database backup and CDS software recovery procedures. This section contains the following sections:

- Performing Backup and Restore on the CDSM Database, page 8-15
- Using the Cisco CDS Software Recovery CD-ROM, page 8-16
- Recovering the System Software, page 8-19
- Recovering a Lost Administrator Password, page 8-21
- Recovering from Missing Disk-Based Software, page 8-22
- Recovering CDS Network Device Registration Information, page 8-25

Performing Backup and Restore on the CDSM Database

The CDSM stores CDS network-wide device configuration information in its Centralized Management System (CMS) database. You can manually back up the CMS embedded database contents for greater system reliability.

To back up the CMS database for the CDSM, use the **cms database backup** EXEC command.

Note The naming convention for backup files includes the timestamp.

To back up and restore the CMS database on the CDSM, do the following:

Step 1 Back up the CMS database to a file.

```
CDE# cms database backup
creating backup file backup-db-11-06-2007-13-10.dump
backup file local1/backup-db-11-06-2007-13-10.dump is ready.
Please use 'copy' commands to move the backup file to a remote host.
```

Step 2 Save the file to a remote server by using the **copy disk ftp** command. This command copies the file from the local disk to a remote FTP server, as shown in the following example:

```
CDE# cd /local1
CDE# copy disk ftp 10.86.32.82 /incoming cds-db-9-22-2002-17-36.dump
cds-db-9-22-2002-17-36.dump

Enter username for remote ftp server:ftp
Enter password for remote ftp server:*******
Initiating FTP upload...
Sending:USER ftp
10.86.32.82 FTP server (Version wu-2.6.1-18) ready.
Password required for ftp.
Sending:PASS *******
User ftp logged in.
Sending:TYPE I
```

```
Type set to I.
Sending:PASV
Entering Passive Mode (10,86,32,82,112,221)
Sending:CWD /incoming
CWD command successful.
Sending PASV
Entering Passive Mode (10,86,32,82,203,135)
Sending:STOR cds-db-9-22-2002-17-36.dump
Opening BINARY mode data connection for cds-db-9-22-2002-17-36.dump.
Transfer complete.
Sent 18155 bytes
```

Step 3 Delete the existing CMS database.

```
CDE# cms database delete
```

Step 4 Restore the CMS database contents from the backup file.

```
CDE# cms database restore cds-db-9-22-2002-17-36
```

Step 5 Enable CMS.

```
CDE# cms enable
```

Using the Cisco CDS Software Recovery CD-ROM

A software recovery CD-ROM image (.iso file) is available for each software release. The recovery CD-ROM can be used to recover system software that must be completely reimaged. The recovery CD-ROM image contains the system software for a single software release and a single application software.

This section presents instructions for creating and using the software recovery CD-ROM to reinstall your system software if for some reason the software that is installed has failed.

Caution If you upgraded your software with a later release than the software recovery CD-ROM image file you downloaded, using the CD-ROM software recovery images may downgrade your system.

System Software Components

Cisco CDS software consists of three basic components:

- Disk-based software
- Flash-based software
- Hardware platform cookie (stored in flash memory)

All of these components must be correctly installed for Cisco CDS software to work properly.

The software is contained in two types of software images provided by Cisco:

- A .bin image containing disk and flash memory components

  An installation containing only the CDS flash memory-based software, without the corresponding disk-based software, boots and operates in a limited mode, allowing for further disk configuration before completing a full installation.

- A .sysimg image containing a flash memory component only

The .sysimg component is provided for recovery purposes, and allows for repair of flash memory only, without modifying the disk contents.

Getting the Cisco CDS Software Recovery File from Cisco.com

To get a software file from Cisco.com, do the following:

Step 1 Launch your web browser and enter the following URL:

http://www.cisco.com/kobayashi/sw-center/sw-video.shtml

The Log In page is displayed.

Step 2 Log in to Cisco.com using your designated username and password. The Video and Content Delivery page is displayed, listing the available software products.

Step 3 Click **Cisco Content Delivery Systems (CDS)**. The Downloads page is displayed.

Step 4 Click the **Cisco Content Delivery Applications** folder to expand it, and click the **Cisco Internet Streamer Application**. The page refreshes and the software releases are displayed.

Step 5 Click the link for the software recovery file you want to download.

- If this is the first time you have downloaded a file from Cisco.com, the Cisco Systems Inc., Encryption Software Usage Handling and Distribution Policy is displayed. Read the policy, fill in the unfilled fields, and click **Accept**.
- If you previously filled out the Encryption Software Usage and Handling and Distribution form, the form does not display again.

The Download page is displayed with the information about the software image file and a Download link.

Step 6 Click **Download**. The Cisco End User Software License Agreement is displayed.

Step 7 Read the agreement and click **Agree**. The File Download dialog box is displayed.

Step 8 Click **Save**. The Save As dialog box is displayed.

Step 9 Navigate to the location where you want to save the file and click **Save**. The file downloads.

Step 10 Burn the software recovery image file onto a CD-ROM.

Installing the Software Using the Recovery CD-ROM

Note For the CDE200, do not use the internal CD-ROM drive for recovery or remanufacture procedures. Use the external USB CD-ROM drive for all procedures that require a CD-ROM.

For the CDE100, use the internal CD-ROM drive for recovery or remanufacture procedures.

To install the system software by using the recovery CD-ROM, perform the following steps:

Step 1 Plug a USB CD-ROM drive into a USB port on the device.

Step 2 Insert the recovery software CD-ROM into the USB CD-ROM drive, and boot the device.

Step 3 When the installer menu appears, choose **Option 7: Wipe Out Disks and Install .bin Image**. (The installer menu options are described in the next section.)

Step 4 Wait for the process to complete.

Step 5 Before you reboot the device, remove the USB CD-ROM drive from the USB port so that the device boots from flash memory.

Step 6 Reboot the device by choosing **Option 8: Exit and Reboot**.

Recovery CD-ROM Options

The options described in the following sections are available from the software recovery CD-ROM installer menu.

Option 1: Configure Network

If the .bin image you need to install is located on the network instead of the CD-ROM, then you must choose this option to configure the network before attempting to install the .bin image.

This option is automatically performed if you install a .sysimg file from the network.

Option 2: Manufacture Flash

This option verifies the flash memory and, if the memory is invalid, automatically reformats it to contain a Cisco standard layout. If reformatting is required, a new cookie is automatically installed.

This option is automatically performed as part of a .bin or .sysimg installation.

Option 3: Install Flash Cookie

This option generates a hardware-specific platform cookie and installs it in flash memory. This option only needs to be performed if there has been a change in the hardware components, such as replacing the motherboard, or moving a flash memory card between systems.

This option is automatically performed during the flash manufacturing process, if needed, as part of a .bin or .sysimg installation.

Option 4: Install Flash Image from Network and Option 5: Install Flash Image from CD-ROM

These options allow installation of the flash memory .sysimg only, and do not modify disk contents. They may be used when a new chassis has been provided and populated with the customer's old disks that need to be preserved.

These options automatically perform flash verification and hardware cookie installation, if required. When installing from the network, you are prompted to configure the network if you have not already done so.

Option 6: Install Flash Image from Disk

This option is reserved for future expansion and is not available.

Option 7: Wipe Out Disks and Install .bin Image

Caution Option 7 erases the content from all disk drives in your device.

This option provides the preferred procedure for installing the Cisco CDS software. This option performs the following steps:

1. Checks that flash memory is formatted to Cisco specifications.

   If yes, continues to number 2.

   If no, the following takes place:

   a. Reformats the flash memory, which installs the Cisco file system.

   b. Generates and installs a platform-specific cookie for the hardware.

2. Erases data from all drives.
3. Remanufactures the default Cisco file system layout on the disk.
4. Installs the flash memory component from the .bin image.
5. Installs the disk component from the .bin image.

Option 8: Exit and Reboot

This option reboots the device. Remove the USB CD-ROM drive before rebooting in order to boot from flash memory.

Recovering the System Software

The Service Engine, Service Router, and CDSM have a resident rescue system image that is invoked should the image in flash memory be corrupted. A corrupted system image can result from a power failure that occurs while a system image is being written to flash memory. The rescue image can download a system image to the main memory of the device and write it to flash memory.

Note The .sysimg file is located under the images folder on the Recovery CD-ROM. If you have upgraded the CDS software, download the corresponding rescue CD iso image, copy to a CD and use the rescue iso image.

To install a new system image using the rescue image, follow these steps:

Step 1 Download the system image file (*.sysimg) to a host that is running an FTP server.

Step 2 Establish a console connection to the device and open a terminal session.

Step 3 Reboot the device by toggling the power switch.

The rescue image dialog appears. The following example demonstrates how to interact with the rescue dialog (user input is denoted by entries in bold typeface):

```
This is the rescue image. The purpose of this software is to let
you download and install a new system image onto your system's
boot flash device. This software has been invoked either manually
(if you entered `***' to the bootloader prompt) or has been
invoked by the bootloader if it discovered that your system image
```

```
in flash had been corrupted.

To download an image, this software will request the following
information from you:
    - which network interface to use
    - IP address and netmask for the selected interface
    - default gateway IP address
    - server IP address
    - which protocol to use to connect to server
    - username/password (if applicable)
    - path to system image on server

Please enter an interface from the following list:
    0: FastEthernet 0/0
    1: FastEthernet 0/1
    0
Using interface FastEthernet 0/0

Please enter the local IP address to use for this interface:
[Enter IP Address]: 172.16.22.22

Please enter the netmask for this interface:
[Enter Netmask]: 255.255.255.224

Please enter the IP address for the default gateway:
[Enter Gateway IP Address]: 172.16.22.1

Please enter the IP address for the FTP server where you wish
to obtain the new system image:
[Enter Server IP Address]: 172.16.10.10

Please enter your username on the FTP server (or 'anonymous'):
[Enter Username on server (e.g. anonymous)]: anonymous

Please enter the password for username 'anonymous' on FTP server (an email address):

Please enter the directory containing the image file on the FTP server:
[Enter Directory on server (e.g. /)]: /

Please enter the file name of the system image file on the FTP server:
[Enter Filename on server]: CDS23.sysimg

Here is the configuration you have entered:
Current config:
          IP Address: 172.16.22.22
             Netmask: 255.255.255.224
     Gateway Address: 172.16.22.1
      Server Address: 172.16.10.10
            Username: anonymous
            Password:
     Image directory: /
      Image filename: CDS-23.sysimg

Attempting download...
Downloaded 34234368 byte image file
A new system image has been downloaded.
You should write it to flash at this time.
Please enter 'yes' below to indicate that this is what you want to do:
[Enter confirmation ('yes' or 'no')]: yes
Ok, writing new image to flash
.....................................................................................Finished
writing image to flash.
Enter 'reboot' to reboot, or 'again' to download and install a new image:
[Enter reboot confirmation ('reboot' or 'again')]: reboot
```

```
Restarting system.
Initializing memory. Please wait.
```

Step 4 Log in to the device as username **admin**. Verify that you are running the correct version by entering the **show version** command.

```
Username: admin
Password:

Console> enable
Console# show version
Content Delivery System Software (CDS)
Copyright (c) 2007 by Cisco Systems, Inc.
Content Delivery System Software Release 2.0.0 (build b460 July 5 2007)
Version: se507-2.0.0

Compiled 02:34:38 July 15 2007 by (cisco)
Compile Time Options: PP SS

System was restarted on Thu July 15 16:03:51 2007.
The system has been up for 4 weeks, 1 day, 6 hours, 7 minutes, 23 seconds.
```

Recovering a Lost Administrator Password

If an administrator password is forgotten, lost, or misconfigured, you will need to reset the password on the device.

Note There is no way to restore a lost administrator password. You must reset the password to a new one, as described in this procedure.

To reset the password, do the following:

Step 1 Establish a console connection to the device and open a terminal session.

Step 2 Reboot the device.

While the device is rebooting, watch for the following prompt and press **Enter** when you see it:

```
Cisco CDS boot:hit RETURN to set boot flags:0009
```

Step 3 When prompted to enter bootflags, enter the **0x800** value.

```
Available boot flags (enter the sum of the desired flags):
0x4000 - bypass nvram config
0x8000 - disable login security

[SE boot - enter bootflags]:0x8000
You have entered boot flags = 0x8000
Boot with these flags? [yes]:yes

[Display output omitted]
Setting the configuration flags to 0x8000 lets you into the system, bypassing all
security. Setting the configuration flags field to 0x4000 lets you bypass the NVRAM
configuration.
```

Step 4 When the device completes the boot sequence, you are prompted to enter the username to access the CLI. Enter the default administrator username (**admin**).

```
Cisco Service Engine Console

Username: admin
```

Step 5 When you see the CLI prompt, set the password for the user using the **username password** command in global configuration mode.

```
ServiceEngine# configure
ServiceEngine(config)# username admin password 0 password
```

You can specify that the password be either clear text or encrypted. The password strength must be a combination of alphabetic character, at least one number, at least one special character, and at least one uppercase character.

Note Do not set the user ID (uid).

Step 6 Save the configuration change by using the **write memory** command in EXEC mode.

```
ServiceEngine(config)# exit
ServiceEngine# write memory
```

Step 7 Optionally, reboot your device by using the **reload** command.

```
ServiceEngine# reload
```

Rebooting is optional; however, you might want to reboot to ensure that the boot flags are reset, and to ensure that subsequent console administrator logins do not bypass the password check.

Note In CDS software, the bootflags are reset to 0x0 on every reboot.

Recovering from Missing Disk-Based Software

This section describes the recovery procedures to use if for some reason the software installation on both system disks is corrupt or missing.

There are two types of disk volumes in the CDS: sysfs and cdnfs. A disk is either allocated as a sysfs disk or a cdnfs disk. The sysfs, or system volumes, contain data and applications that are critical to the system's basic functionality.

The sysfs volumes are stored in a two-disk RAID-1 (mirrored) array. RAID-1 duplicates data between each of the disks in the array. The two-disk scheme allows for either of the drives in the sysfs volumes array to fail without sustaining data loss or incurring system errors.

The status of the volumes can be seen through the **show disk raid-state** command, and can be in any of the following states:

- Normal—Both drives are attached, and data is mirrored between them.
- Syncing—Data is being copied between the drives to restore the volumes to a normal state. This typically happens when a new drive is added to repair degraded volumes.

- Degraded—One of the disks has failed. It is highly recommended that a new disk is added to repair the volumes.
- Bad—Both disks have failed. The system has likely lost all but basic functionality.

Note If both system disks fail, a CDS state of "missing disk-based software" occurs. Normally, when a problem occurs on one system disk, a disk failure or RAID alarm is triggered. If this occurs, replace the failed disk. See the "Disk Maintenance" section on page 8-26.

The CDS state of "missing disk-based software" is most likely to occur if you replaced both system disks in your Service Engine, Service Router, or CDSM. By design, the software installation on the system disks cannot be corrupted by a system failure or a power failure.

If both system disks fail or are missing, the software continues to run. However, it runs in a basic functionality mode in which HTTP proxy and related HTTP features still work, but most other features fail.

Caution This procedure should only be used as a last-resort method to recover the system software on a unit. Typically, the sysfs automatically repairs itself across a reboot if any new disks are detected. If the volumes are degraded and a new disk is present at reboot, the new disk is added to the existing array (sync starts). If the volumes are "bad" and a new disk is present at reboot, the initial sysfs volume will be built on the disk.

To recover from this condition, do the following:

Step 1 Remove the Service Engine record from the CDSM.

a. Choose **Devices > Devices**.

b. Click the **Edit** icon next to the name of the Service Engine that you want to delete. The Device home page is displayed.

c. Click the **Trash** icon. You are prompted to confirm your decision.

d. Click **OK** to execute your request. The Service Engine is removed from the CDSM.

Note The Service Engine registration record needs to be deleted from the CDSM in order for the Service Engine to complete reregistration after it comes back online. The CDSM will not register a device if the device already appears in the record as registered.

Step 2 Power down the device and replace the failed or missing system disks with new, blank disks.

Step 3 After the new disks are installed, power up the device.

Step 4 From a console or through an SSH session, check the startup messages that appear on your screen.

If there is a problem with the system disk or the disk-based software, a message similar to the following appears:

```
Jan 21 21:55:45 (none) ruby_disk:%SE-DISK-2-200024:First disk not in standard
configuration. Run 'disk recover-system-volumes' command and re-install software.
ruby_disk:Your first disk is not in standard configuration.
ruby_disk:Run 'disk recover-system-volumes' from the CLI

****************************************
    System software is missing.
```

```
    Check whether first-disk is bad, or
    use 'disk recover-system-volumes' to recover first-disk.
*****************************************
```

Step 5 Log in as **admin**.

```
Cisco Service Engine Console

Username: admin
Password:
System Initialization Finished.

SE-507 con now available

Press RETURN to get started!
```

Step 6 After logging in to a console or SSH session, enter the **copy ftp install** or **copy http install** EXEC command to download and install a new system image.

```
ServiceEngine# copy ftp install ftp-server remotefiledir remotefilename
```

For example:

```
SE# copy ftp install vista /CDS/upgrades CDS-2.0.0.2-K9.bin
Enter username for remote ftp server: biff
Enter password for remote ftp server:
Initiating FTP download...
printing one # per 1MB downloaded
Reclaiming unused safe state sectors...
################################################################################
############
################################################################################
############
Installing phase3 bootloader...
Installing system image to flash: done
The new software will run after you reload.
#
ServiceEngine# show flash
CDS software version (disk-based code): CDS-2.0.1-b130

System image on flash:
Version: 2.0.1.2

System flash directory:
System image: 98 sectors
Bootloader, rescue image, and other reserved areas: 26 sectors
128 sectors total, 4 sectors free.
```

Step 7 Reboot the software with the new disk and new system image by entering the **reload** EXEC command.

```
SE# reload
```

Step 8 Register the device with the CDSM by using the **cms enable** command in global configuration mode.

```
SE# configure
SE(config)# cms enable
```

Recovering CDS Network Device Registration Information

Device registration information is stored both on the device itself and on the CDSM. If a device loses its registration identity or needs to be replaced because of hardware failure, the CDS network administrator can issue a CLI command to recover the lost information or, in the case of adding a new device, assume the identity of the failed device.

To recover lost registration information, or to replace a failed node with a new one having the same registration information, do the following:

Step 1 Mark the failed device as "Inactive" and "Replaceable" in the CDSM.

a. Choose **Devices > Devices**.

b. Click the **Edit** icon next to the name of the Service Engine you want to deactivate. The Device home page is displayed.

c. From the left-panel menu, choose **Device Activation**.

d. Uncheck the **Activate** check box. The page refreshes, displaying a check box for marking the device as replaceable.

e. Check the **Replaceable** check box and click **Submit**.

Note This check box only displays when the device is inactive.

Step 2 Configure a system device recovery key.

a. Choose **System > Configuration**.

b. Click the **Edit** icon next to the System.device.recovery.key property. The Modifying Config Property page is displayed.

c. Enter a password in the **Value** field and click **Submit**. The default password is **default**.

Step 3 Configure the basic network settings for the new device.

Step 4 Open an SSH session to the device CLI and enter the **cms recover identity** *keyword* EXEC command, where *keyword* is the device recovery key that you configured in the CDSM.

When the CDSM receives the recovery request from the Service Engine, it searches its database for the Service Engine record that meets the following criteria:

- The record is inactive and replaceable.
- The record has the same hostname as given in the recovery request.
- The device is the same hardware model as the device in the existing record.
- The file system allocations for the device are the same as or greater than the device in the existing record.

If the recovery request matches the Service Engine record, then the CDSM updates the existing record and sends the requesting Service Engine a registration response. The replaceable state is cleared so that no other device can assume the same identity. When the Service Engine receives its recovered registration information, it writes it to file, initializes its database tables, and starts.

Step 5 Return to the CDSM and activate the device.

a. Choose **Devices > Devices**.

b. Click the **Edit** icon next to the name of the Service Engine you want to activate. The Device home page is displayed.

c. From the left-panel menu, choose **Device Activation**. The Service Engine status should be Online.

d. Check the **Activate** check box and click **Submit**.

Disk Maintenance

This section covers removing and replacing disk drives, and manual content removal.

For information about disk error handling, see the "Enabling Disk Error Handling" section on page 4-56.

Removing and Replacing Hard Disk Drives

In brief, the procedure for replacing a disk is simply to power down the unit, remove the disk, insert the new disk, and reboot. During the reboot, the system automatically detects any new disks and seamlessly allocates their space according to a simple disk-policy heuristic.

The disk policy's design, when adding new disks, is to always favor safety. If when a new disk is added, the disk manager detects "degraded" or "bad" sysfs volumes, the new disk will be used to repair the sysfs volumes. Thus, the disk manager will always strive to have two disks allocated to the sysfs volumes. If when a new disk is added, the sysfs volumes are "normal" or "syncing," the new disk is added to the cdnfs volume.

See Figure 8-2 for the disk numbering assignment.

Figure 8-2 Disk Numbering on a CDE200

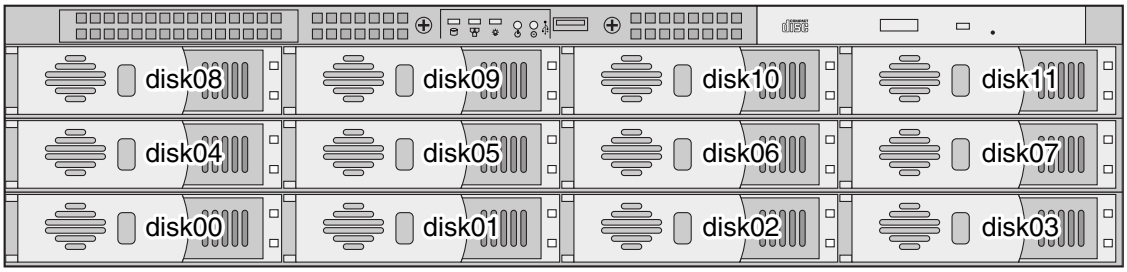

Specifically, the procedure for removing a disk is different for the two types of disk volumes (sysfs and cdnfs). The two procedures are outlined in the following sections.

Removing and Replacing a cdnfs Disk

To remove and replace a cdnfs disk, do the following:

Step 1 Stop all applications and unmount the drive by using the **disk unuse** command. This command stops and restarts all applications that are currently using the specified disk drive (for example, disk02 [/local/local2] or disk03). All file systems, if applicable, will be unmounted on the specified disk.

The syntax of the command is as follows:

disk unuse *diskname*

The variable *diskname* is the name of the drive to be unmounted. See Figure 8-2 for the disk numbering used for the disk name.

Following are some examples of the command and the resultant actions:

```
se# disk unuse disk01
Disk01 has mounted SYSFS and can not be unused!

se# disk unuse disk02
This will restart applications currently using disk02 and unmount all partitions on
disk02.
Do you want to continue? (yes/no) [no]no
Disk02 not unused.

se# disk unuse disk02
This will restart applications currently using disk02 and unmount all partitions on
disk02.
Do you want to continue? (yes/no) [no]yes
Disk02 has been unused. No application is using disk02 now.

se# disk unuse disk02 delete-partitions
This will restart applications currently using disk02 and unmount and *delete* all
partitions on disk02.
Do you want to continue? (yes/no) [no]yes
Disk02 has been unused. No application is using disk02 now.
And all partitions on disk02 are deleted.
```

Step 2 Power down the unit.

Step 3 Remove the bad disk and insert the new disk.

Step 4 Power up the unit.

Removing and Replacing a sysfs Disk

For sysfs disk drives, you need to mark the disk as bad before powering down the unit. Then, after replacing the disk and powering up the unit, mark the disk as good and reboot.

To remove and replace a sysfs disk, do the following:

Step 1 Mark the sysfs disk as bad.

```
SE# disk mark disk03 bad
disk03 is marked as bad.
It will not be used after reload.
```

Step 2 Reboot the unit by entering the **reload** EXEC command.

```
SE# reload
```

Step 3 Replace the disk.

Step 4 Power up the unit.

Step 5 After logging in to a console or SSH session, mark the disk as good.

```
se# disk mark disk03 good
disk03 is marked as good.
It will be used after reload.
```

Step 6 Reboot the unit by entering the **reload** EXEC command.

```
SE# reload
```

Removing Content

The Content Removal page allows you to delete content manually, either by referencing a single URL or by using an XML file for batch deletions.

Single Content Removal

To delete content manually, do the following:

Step 1 Choose **Services > Service Definition > Delivery Service**. The Delivery Services Table page is displayed

Step 2 Click the **Edit** icon next to the delivery service name that has the content you want to delete.

Step 3 Choose **Tools > Content Management**. The Content Removal page is displayed.

Step 4 In the **URL String** field for a single URL removal, enter the URL of the content you want to remove.

Click **Help** for an example.

Step 5 Click **Submit**.

Batch Content Removal

The batch content removal uses an XML file that lists the URLs of the content to delete. The XML batch content removal file can be created with any ASCII text editing tool.

The XML Schema file describes and dictates the content of the XML file. The URLRemoval.xsd file contains the XML schema.

The following code is the URL Removal XML schema:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
            xmlns:cds-ns="http://cisco.com/unicorn/cds/urlmgmt"
            targetNamespace="http://cisco.com/unicorn/cds/urlmgmt"
            xmlns="http://cisco.com/unicorn/cds/urlmgmt"
            elementFormDefault="qualified">

    <xsd:element name="URLRemovalList">
        <xsd:complexType>
            <xsd:sequence>
                <xsd:element name="url-entry" minOccurs="1" maxOccurs="unbounded" type="xsd:string" />
            </xsd:sequence>
        </xsd:complexType>
    </xsd:element>

</xsd:schema>
```

To delete content items in a batch, do the following:

Step 1 Create an XML file containing the URLs of the content you want to delete by using the XML example shown in the Content Removal—Batch File Help page.

Figure 8-3 shows the URL Removal XML file example displayed when you click **Help**.

Figure 8-3 Content Removal—Batch File Help Page

Step 2 In the **Batch File** field for a Batch URL removal, enter the URL of the XML file.

Step 3 Click **Submit**.

APPENDIX A

Troubleshooting

This appendix provides information on troubleshooting. The following topics are covered in this appendix:

- Troubleshooting Service Router Configurations, page A-1
- Troubleshooting the Distribution Hierarchy, page A-2
- Troubleshooting Content Acquisition, page A-3
- Enabling the Kernel Debugger, page A-6

For more troubleshooting tools, see Chapter 7, "Monitoring the Internet Streamer CDS."

Troubleshooting Service Router Configurations

Because there are many steps required for the Service Router to redirect the request properly, you might see some content request errors from the Service Router when the configuration is not quite complete. Here are some areas to look at when troubleshooting:

- DNS delegation
  - Is the requested domain delegated to the Service Router on the DNS server that is authoritative for the parent domains? The Service Router's DNS name should be forward resolvable. Check with the system administrator to delegate a domain.
- Service Router routing properties
  - Is the Service Router activated? See the "Activating a Service Router" section on page 4-87 to activate a Service Router.
  - Is a default coverage zone set for a Service Engine, or is there a CDS network-wide Coverage Zone file or a local Coverage Zone file set for the Service Router? See the "Coverage Zone File Registration" section on page 6-9 to set a Coverage Zone file. See the "Configuring the Service Engine" section on page 4-6 to set a default coverage zone.
  - Is the content request from an end system covered by a Service Engine in a coverage zone based on the default coverage zone or the Coverage Zone file? This Service Engine is the "serving Service Engine." See the "Coverage Zone File" section on page 1-21 for information on coverage zones. See Appendix C, "Creating Coverage Zone Files," for information on creating a Coverage Zone file.
  - Is the serving Service Engine activated? See the "Activating a Service Engine" section on page 4-6 to activate a Service Engine.

  - Is there a delivery service created for the requested domain and a serving Service Engine assigned to this delivery service? See the "Creating Delivery Service" section on page 5-4.
  - Is the serving Service Engine alive? Use the **show statistics service-routing se** command to show the status of a Service Engine. See the "Using show and clear Commands" section on page 7-9.
- Content prefetched on a Service Engine
  - Is a Manifest file assigned to the delivery service associated with the serving Service Engine? See the "Working with Manifest Files" section on page B-2.
  - Is the Manifest file accessible from the CDSM? See the "Identifying Content Using a Manifest File" section on page 5-22.
  - Is there any syntax error in the Manifest file? See the "Manifest File Structure and Syntax" section on page B-17.
  - Is the requested content specified in the Manifest file? See the "Specifying a Single Content Item" section on page B-2.
  - If the requested content is streaming media, is the protocol engine enabled? See the "Application Control" section on page 4-28.

Troubleshooting the Distribution Hierarchy

Because distribution-related problems are design-dependent, your initial strategy is to discover whether or not the correct Service Engine is sending content in the correct distribution path.

- To determine which Service Engines are in the distribution path of a particular Service Engine, use the **show distribution remote traceroute** EXEC command, as shown in the following example:

```
ce1# show distribution remote traceroute ?
  forwarder-next-hop  next forwarder along the path
  unicast-sender      check status for unicast sender

ce1# show distribution remote traceroute forwarder-next-hop ?
  delivery-service-id  Delivery-service-id of a Delivery Service

ce1# show distribution remote traceroute forwarder-next-hop delivery-service-id 133 ?
  max-hop          Trace route till specified number of hops is reached
  trace-till-good  traceroute till probe is good or the object is found
  trace-till-root  traceroute till the acquirer

ce1# show distribution remote traceroute forwarder-next-hop delivery-service-id 133
trace-till-root

Hop NextHop_SEId  NextHop_SEName NextHop_SEIp      GenID Status/Reason
--- ------------  -------------- ------------      ----- -------------
  1    1100               ce3     10.255.0.43     1       LOC-LEAD
  1    1100               ce3 128.107.193.183     1       LOC-LEAD (Reached RootCE)
```

- To verify that the Service Engine is reachable and that it is in the distribution hierarchy, use the **show distribution remote traceroute** EXEC command, as shown in the following example:

```
se1# show distribution remote traceroute unicast-sender delivery-service-id 133 ?
  cdn-url           check the object on remote SE using cdn-url
  probe             probe the remote unicast sender
  relative-cdn-url  check the object on remote SE using relative-cdn-url
```

```
se1# show distribution remote traceroute unicast-sender delivery-service-id 133 probe
?
  max-hop          Max-hop to traceroute to
  trace-till-good  traceroute till probe is good or the object is found
  trace-till-root  traceroute till the root se

se1# show distribution remote traceroute unicast-sender delivery-service-id 133 probe
trace-till-root
Polling .... se3 [10.255.0.43]  Fwdr_Id:1100
Polling .... se3 [128.107.193.183]  Fwdr_Id:1100
(Reached RootSE)
```

Troubleshooting Content Acquisition

To monitor acquisition progress and to troubleshoot, use the following commands from the Content Acquirer CLI:

- Use the **show acquirer delivery-services** EXEC command to obtain delivery service information, such as the delivery service ID and delivery service name, that you need to enter in other **show acquirer** commands, such as the **show acquirer progress** command. In the following example, the delivery service ID is 793 and the delivery service name is group01-cifs.

```
SE# show acquirer delivery-services
Querying Database.......
Acquirer information for all delivery services:
---------------------------------------
Delivery-service-id       : 793
Delivery-service-Name     : group01-cifs
WebSite-Name     : group01-cifs
Root-CE-Type     : Configured
State            : Enabled
Disk Quota       : 200 MB
Origin FQDN      : cdn.allcisco.com
Delivery-service Priority : 500
Manifestfile-TTL : 5
Manifestfile-URL : ftp://10.1.1.1/cifs.xml
```

- Use the **show acquirer** EXEC command to make sure that the acquirer process on the Content Acquirer is working correctly, and that the device is using the expected amount of bandwidth for acquisition. The following example shows that the acquirer is running properly and that the device is configured with unlimited bandwidth for acquisition of content.

```
SE# show acquirer
Acquirer is running OK
Current Acquisition Bandwidth:Not Limited
```

- Use the **show acquirer progress** EXEC command to check how far the acquisition of content has progressed. A specific delivery service ID or delivery service name can be specified to obtain the progress for a specific delivery service. In the following example, the acquirer has already acquired 2237 items.

```
SE# show acquirer progress delivery-service-id 793
Querying Database.......

Acquirer progress information for delivery service ID:793
Delivery-service-Name:group01-cifs
-----------------------------------------------------------------

Acquired Single Items   :         0 / 0
Acquired Crawl Items    :     2237 / 2500     -- start-url=www.mtv.com//
```

- Use the **show statistics acquirer delivery-service-id** or **show statistics acquirer delivery-service-name** EXEC command to obtain the detailed acquisition statistics for a given delivery service. In the following example, there was an error acquiring two items.

```
SE# show statistics acquirer delivery-service-id 793
Querying Database.......

Statistics for Delivery Service Delivery-service-id :793 Delivery-service-Name
:group01-cifs
-------------------------------------------------------

Manifest:
---------
        Fetch Errors    :0
        Parsing Errors  :0
        Parsing Warnings:0

Acquisition:
------------
        Total Number of Acquired Objects        :2237
        Total Disk Used for Acquired Objects    :981511280 Bytes
        Total Number of Failed Objects          :2
        Total Number of Re-Check Failed Objects :0
```

- Use the **show statistics acquirer errors delivery-service-id** or **show statistics acquirer errors delivery-service-name** EXEC command to see the reasons why the errors occurred. In the following example, one error occurred because there was a problem acquiring the URL. The other error occurred because the disk quota for the delivery service configured in the Content Distribution Manager GUI would have been exceeded if the specified URL had been acquired. You can increase the delivery service disk quota to correct this error.

```
SE# show statistics acquirer errors delivery-service-id 793
Querying Database.......

Acquisition Errors for the Delivery Service ID:793
-------------------------------------
Crawl job:start-url http://www.mtv.com//
Crawl Errors
---------------
Internal Server Error(500):http://cgi.cnn.com/entries/intl-emailsubs-confirm
Exceeded Disk Quota(703):http://www.cdt.org/copyright/backgroundchart.pdf
```

- If more detailed troubleshooting of content acquisition is required, you can increase the debug level of the acquirer using the **debug acquirer trace** EXEC command. The logs are written to local1/errorlog/acquirer-errorlog.current.
- To verify that an expected object has been pre-positioned on the Service Engine, use the **show distribution object-status** EXEC command, as shown in the following example:

```
SE# show distribution object-status
http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv

        ========== Website Information ==========
Name:           RTPServer5
Origin Server FQDN:       172.18.81.168
Service Routing FQDN:     N/A
Content UNS Reference #:  1
        ========== Delivery Services Information ==========
        *** Delivery Service 1903 (name = A_Multicast) ***

Object Replication
------------------
```

```
Replication:              Done
File State:               Ready for distribution
Multicast for Delivery Service: Not Enabled
Replication Lock:         Received by Unicast-Receiver/Acquirer
Reference Count:          1
Total Size:               2756437
Transferred Size:         2756437
MD5 of MD5:               tjS#DxqE5oUc024Z8XtFDw..
Source Url:               http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv
Source Last Modified Time:Wed Jan  7 19:03:48 2004

Object Properties
-----------------
Redirect To Origin:       Yes
Requires Authentication:  No
Alternative URL:
Serve Start Time:         N/A
Serve End Time:           N/A
Play servers:             HTTP HTTPS WMT
Content Metadata:         None
Content uns_id:           NgcJTCU#JaY4ZGIPbsrONw..
Content gen-id:           1768:1136512329:2

========== CDNFS Information ==========
Internal File Name:
/disk00-04/d/http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv
Actual File Size:         2756437 bytes
MD5 of MD5 (Re-calculated):tjS#DxqE5oUc024Z8XtFDw..
Content metadata:         None
Metadata match with:      Delivery Service 1903
Number of Source-urls:    1

Source-url to CDN-object mapping:
Source-url:               http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv
Used by CDN object:       ---- Yes ----
Internal File Name:
/disk00-04/d/http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv
Actual File Size:         2756437 bytes

 ========== CDNFS lookup output ==========
CDNFS File Attributes:
  Status                  3  (Ready)
  File Size               2756437 Bytes
  Start Time              null
  End Time                null
  Allowed Playback via    HTTP WMT HTTPS
  Last-modified Time      Wed Jan  7 19:03:48 2004
  cache-control           max-age=864000
  cdn_uns_id              NgcJTCU#JaY4ZGIPbsrONw..
  content-type            video/x-ms-wmv
  etag                    "042e6fa50d5c31:b39"
  file_duration           65
  last-modified           Wed, 07 Jan 2004 19:03:48 GMT
  server                  Microsoft-IIS/6.0
  x-powered-by            ASP.NET
Internal path to data file:
/disk00-04/d/http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv
```

By comparing fields, such as Total Size, Transferred Size, and Source URL in the Object Replication output and Actual File Size and Source URL in the Source-URL to CDN-Object Mapping output, you can determine whether or not the object that is stored is the same as the object that was requested.

- To view the file directory structure on the Service Engine and verify the physical file on the disk, you can use the **cdnfs browse** EXEC command, as shown in the following example:

```
SE# cdnfs browse

------  CDNFS interactive browsing  ------
dir, ls:   list directory contents
cd,chdir:  change current working directory
info:      display attributes of a file
more:      page through a file
cat:       display a file
exit,quit: quit CDNFS browse shell

/>ls
                    srv2-static.cisco.com/
                    http-172.18.81.151-glfc4h-b9gywnf5rlnfweg/
                    http-172.18.81.163-og5o21u178nrhw1mctgtiq/
                    172.18.81.163/
                    file--xrnfwxifgu62jtiwtyixvg/
                    172.18.81.168/
                    http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/
                    172.18.81.155/
/>cd 172.18.81.168/
/172.18.81.168/>ls
376 Bytes           manifest-Delivery_service_1903.xml-lEYmrfnjt2o5GbUNwLCApA
                    Videos/
/172.18.81.168/>cd Videos
/172.18.81.168/Videos/>ls
2756437 Bytes       SM-final Innebandy 2003.wmv <===============Physical file on disk
/172.18.81.168/Videos/>quit
SE#
```

Enabling the Kernel Debugger

Cisco CDS software allows you to enable or disable access to the kernel debugger from the CDSM. Once enabled, the kernel debugger is automatically activated when kernel problems occur.

To enable the kernel debugger, do the following:

Step 1 Choose **Devices > Devices > General Settings > Troubleshooting > Kernel Debugger**. The Kernel Debugger page is displayed.

Step 2 To enable the kernel debugger, check the **Enable** check box, and click **Submit**.

APPENDIX B

Creating Manifest Files

This appendix describes the process for creating Manifest files used to acquire and distribute content within the CDS network. This appendix includes the following topics:

- Introduction, page B-1
- Working with Manifest Files, page B-2
- Manifest Validator Utility, page B-14
- Manifest File Structure and Syntax, page B-17
- XML Schema, page B-44
- Manifest File Time Zone Tables, page B-50

For information about using a Manifest file in a delivery service, see the ""Identifying Content Using a Manifest File" section on page 5-22.

Introduction

The CDS is used to ingest, distribute, and deliver multi-format content to different client devices. In order to specify the content to be prefetched and to control the delivery of the prefetched content, an XML file called a *Manifest file* is used. Third-party asset management systems can interoperate with the CDS by using this Manifest file interface. Each delivery service in the CDS can be configured with or without a Manifest file. The Manifest file can also be automatically generated by using the CDSM. The Manifest file is primarily used in prefetch ingest and hybrid ingest.

The Manifest file is specified in the CDSM in the following ways:

- External Manifest File Specification—The Manifest file is hosted on an external server and a URL pointing to that server is configured in the delivery service. The Manifest file can be fetched using FTP, HTTP, HTTPS and CIFS protocols.
- GUI Configured—The CDSM GUI can generate a Manifest file. The CDSM provides the required elements for the user to create a Manifest file and to specify the attributes in the Manifest file. Only commonly used attributes are supported by the CDSM.

The Manifest file is processed by the Content Acquirer. The Content Acquirer parses the Manifest file, creates the metadata based on the attributes in the file, and prefetches the content specified. For live content and content that is ingested on demand, the Content Acquirer creates the metadata and does not fetch the actual content. The metadata created by the Content Acquirer is propagated to all the Service Engines participating in the delivery service.

Manifest File Requirements

The Manifest file needs to support different attributes and tags to support content prefetching and hybrid ingest. The basic requirements for a Manifest file are the following:

- Specify Content to Be Prefetched—There are two ways to specify prefetched content. One is to use a single item, where users specify a single URL and the Content Acquirer ingests only the content pointed to by this URL. Another way is by using a crawler item, where users specify a crawl job with parameters like start-url, depth, prefix, and reject or accept. In this case, the Content Acquirer crawls the origin server to fetch content based on the parameters.
- Specify Schedule Information—To instruct the Content Acquirer when to ingest the content and how often to check the server for updates.
- Specify Publish Information—Information regarding how content is accessed by the end users; for example, the playserver attribute specifies which server to use for playing the content, the cdn-url attribute specifies which URL is used by end-users to access the content, the *serveStartTime* and *serveStopTime* attributes instruct the CDS when it can serve the content and provides additional metadata for playing.
- Specify Live Streaming Content—The Manifest file can also be used to specify live stream splitting.
- Specify Metadata for Hybrid Ingest Content—For hybrid ingest, the Manifest file can be used to specify the content serve start and stop time for content ingested on demand.

Working with Manifest Files

This section provides Manifest file samples for carrying out specific tasks. Each sample has an associated explanation of its purpose and function. The Manifest file can specify a single content object, a website crawler job, or an FTP server crawler job to acquire prefetched content or to acquire information about live content that is distributed to edge Service Engines later.

Specifying a Single Content Item

Use the <item> tag to specify a single content item, object, or URL. The required *src* attribute is used to specify the relative path portion of the URL. If the server *name* attribute is omitted, the server *name* attribute in the last specified <server> tag above it is used. If there are no <server> tags close by in the Manifest file, the server that hosts the Manifest file is used, which means that the relative URL is relative to the Manifest file URL.

The following example provides an example of a Manifest file that specifies single content items:

```
<CdnManifest>
    <item src="http://www.my-server/test.html" />
    <item src="test.html" />
    <server name="my-origin-server-one">
        <host name="http://www.my-server-one.com/eng/" />
    </server>
    <server name="my-origin-server-two">
        <host name="http://www.my-server-two.com/eng/" />
    </server>
    <item src="project-two.html" />
    <item server="my-origin-server-one" src="project-one.html" />
</CdnManifest>
```

For a single item, you specify the item's URL in the *src* attribute. There are two ways to specify the item URL:

- Specify the *src* attribute with the absolute URL as shown in the following format:

  ```
  proto://username:password@/domain-name:port/file-path/file-name
  ```

  In the example, the first <item> tag uses the full path.

- Specify the origin server information using the <server><host> tags and use the *src* attribute to specify only the relative path.

  In the example, every <item> tag except the first one uses a relative path. The second <item> tag uses the Manifest file server, where test.html is relative to the Manifest file URL. The second <item> tag, "project-two.html," uses "my-origin-server-two." The third <item> tag, "project-one.html," uses "my-origin-server-one."

Specifying a Crawl Job

The crawler feature methodically and automatically searches acceptable websites and makes a copy of the visited pages for later processing. The crawler starts with a list of URLs to visit and identifies every web link in the page, adding these links to the list of URLs to visit. The process ends after one or more of the following conditions are met:

- Links have been followed to a specified depth.
- The maximum number of objects has been acquired.
- The maximum content size has been acquired.

By crawling a site at regular intervals using the time-to-live (or *ttl)* attribute, these links and their associated content can be updated regularly to keep the content fresh. Use the <crawler> tag to specify the website or FTP server crawler attributes. Table B-1 lists the attributes, states whether these attributes are required or optional, and describes their functions.

Table B-1 Website or FTP Server Crawl Job Attributes

Attribute	Description
start-url	(Required) Identifies the URL to start the crawl job from. It can be a full path or a relative path. If it is a relative path, the <server><host> tags are required to specify the origin server information.
depth	(Optional) Defines the level of depth to crawl the specified website. The depth is defined as the level of a website's URL links or FTP server's directory, where 0 is the URL or directory from which the crawl job starts. 0 = Acquire only the starting URL. 1, 2, 3, ... = Acquire the starting URL and its referred files to the depth specified. –1 = Infinite or no depth restriction. If the depth is not specified, the default is used. The default is 20. **Note** It is not advisable to specify a depth of –1 because it will take a long time to crawl a large website and is wasteful if all of the content on that particular website is not required.

Table B-1 Website or FTP Server Crawl Job Attributes (continued)

Attribute	Description

prefix	(Optional) Combines the hostname from the <server> tag and this field to create a full prefix. Only content with URLs that match the full prefix are acquired, as shown in this example:
	`<server name="xx"> <host name="www.cisco.com" proto="https" port=433/> </server>`
	with the following <crawler> tag:
	`prefix="marketing/eng/"`
	The full prefix is "https://www.cisco.com:433/marketing/eng/." Only URLs that match this prefix are crawled. If a web page refers to ".../marketing/ops," the marketing/ops page and its children are not acquired.
	If the prefix is omitted, the crawler checks the default full prefix, which is the hostname portion of the URL from the server. In the example, the default full prefix is "https://www.cisco.com:433."

accept	(Optional) Uses a regular expression to define acceptable URLs to crawl, in addition to having acceptable URLs match a prefix. For example, accept="stock" means that only URLs that meet two conditions are crawled: the URL matches the prefix and also contains the regular expression string "stock."
reject	(Optional) Uses a regular expression to reject a URL if it matches the expression. The URL is first checked for a possible prefix match and then checked for a reject regular expression. If a URL does not match the prefix, it is immediately rejected. If a URL matches both the prefix and the reject regular expression, it is rejected by the expression.
max-number	(Optional) Specifies the maximum number of crawl job objects that can be acquired.
maxTotalSizeInMB *maxTotalSizeInKB* *maxTotalSizeInB*	(Optional) Specifies the maximum size of content that this crawl job can acquire. The size can be expressed in bytes (B), kilobytes (KB), or megabytes (MB). **Note** The maximum size of the file that is acquired is going to be less than the amount of disk space required to store the file. Files, when stored, contain overhead that contributes to the amount of disk space used for the delivery service. This overhead is approximately 20 KB per file. File size and storage overhead need to be taken into account when you are configuring the delivery service disk quota. This attribute replaces the *max-size-in-B/KB/MB* attribute. The *max-size-in-B/KB/MB* attribute continues to be supported for backward compatibility only.
externalPrefixes	(Optional) Specifies additional prefixes for crawl jobs to crawl multiple protocols or multiple websites. Prefixes are separated with a bar (|).
externalServers	(Optional) Specifies additional hosts for crawl jobs. Can be used for multiple host crawl jobs where each host has a different user account. This attribute can be used to refer to the <host> tag with the proper authentication information.

Note If you specify both the *max-number* and *maxTotalSizeIn* attributes as the criteria to use to stop a crawl job, the condition that is met first takes precedence. The crawl job stops either when the maximum number of objects is acquired or when the maximum content size is reached, whichever occurs first. For example, if the crawl job has acquired the maximum number of objects specified in the Manifest file but has not yet reached the maximum content size, the crawl job stops.

The following is an example of a website crawl job:

```
<server name="cisco">
    <host name="http://www.cisco.com/jobs/" />
</server>
<crawler
    server="cisco"
    start-url="eng/index.html"
    depth="10"
    prefix="eng/"
    reject="\.pl"
    maxTotalSizeIn-MB="200"
/>
```

This website crawl job example contains the following attributes:

- The *start-url* path is http://www.cisco.com/jobs/eng/index.html.
- Search to a website link *depth* of 10.
- Search URLs with the *prefix* http://www.cisco.com/jobs/eng/.
- Reject URLs containing .pl (Perl script pages).
- Only crawl until 200 megabytes in total content size are acquired.

If the server *name* attribute is omitted, the server *name* in the last specified <server> tag above it is used. If there are no <server> tags close by in the Manifest file, the server that hosts the Manifest file is used, which means that the relative URL is relative to the Manifest file URL.

Writing Common Regular Expressions

A regular expression is a formula for matching strings that follow a recognizable pattern. The following special characters have special meanings in regular expressions:

. * \ ? [ ] ^ $

If the regular expression string does not include any of these special characters, then only an exact match satisfies the search. For example, “stock” must match the exact substring “stock.”

Scheduling Content Acquisition

Two attributes, *ttl* and *prefetch*, are used to schedule content acquisition. Use *ttl* to specify the frequency of checking the content for freshness, in minutes. For example, to check for page freshness every day, enter ttl=“1440.”

In the following example, page freshness is scheduled to be checked once a day:

```
<item
    src="index.html"
    ttl="1440"
/>
```

In the following example, page freshness is scheduled to be crawled and checked every hour to a link *depth* value of 2:

```
<crawler
    start-url="index.html"
    depth="2"
    ttl="60"
/>
```

If the content is not yet available at a particular URL, the *prefetch* attribute can be used to specify the start time for acquisition at the specified URL. For example, prefetch="2002-06-28 18:35:21" means the content acquisition job can only start on June 28, 2002 and at the specified time.

The following example schedules a crawl of this website every hour to a link *depth* value of 2 to start on November 9, 2001 at 8:45 a.m.

```
<crawler
    start-url="index.html"
    depth="2"
    prefetch="2001-11-09 08:45:12"
    ttl="60"
/>
```

Specifying Shared Attributes

Attributes in single <item> tags can be shared or have the same attribute values. Instead of writing these attributes individually for every <item> tag, you can extract them and place them in a higher-level tag called <item-group>, where these attributes can be shared from this higher-level tag. You can create an <item-group> tag at a level below the <CdnManifest> tag, and write <item> tags into it as subtags, moving shared attributes into the <item-group> tag, as shown in the following example:

```
<?xml version="1.0"?>
<CdnManifest>

<server name="cisco-cco">
  <host name="http://www.cisco.com"
        proto="http" />
</server>

<item-group
    server="cisco-cco"
    ttl="1440"
    type="prepos" >

   <item src="jobs/index.html"/>
   <item src="jobs/index1.html"/>
   <item src="jobs/index2.html"/>
   <item src="jobs/index3.html"/>
   <item src="jobs/index4.html"/>
   <item src="jobs/index5.html"/>

</item-group>

</CdnManifest>
```

You can also use the <options> tag to share attributes at the top-most level of the Manifest file. Shared attributes in the <options> tag can be shared by every <item> tag or by the <crawler> tag in the Manifest file. However, if a shared attribute is specified in both the <item-group> and the <item> tags or the <options> and <item> tags, attribute values in the <item> tags take precedence over the <item-group> and <options> tags.

The following example illustrates this precedence rule. The first <item> tag takes the *ttl* value 1440 from the <options> tag, but the second <item> uses its own *ttl* value of 60.

```
<options
   ttl="1440" >
<item src="index.html" />
<item src="index1.html" ttl="60" />
```

Specifying a Crawler Filter

With a rule-based crawler filter, you can crawl an entire website and only acquire contents with certain predefined characteristics. In contrast, crawler attributes in the <crawler> tag do not act as filters but only define the attributes for crawling. The <matchRule> tag is designed to act as a rule-based filter. You can define rule-based matches for file extensions, size, content type, and timestamp. In the following example, the crawl job is instructed to crawl the entire website starting at "index.html," but to acquire only files with the .jpg extension and those larger than 50 kilobytes.

```
<crawler
   start-url="index.html" >
   <matchRule>
       <match minFileSizeIn-KB="50" extension="jpg" />
    </matchRule>
</crawler>
```

There can be multiple <match> subtags within a <matchRule> tag. Table B-2 lists and describes the <match> subtag attributes.

Table B-2 <match> Subtag Attributes

Attribute	Description
mime-type	Specifies match of these MIME-types.
extension	Specifies match of files with these extensions.
time-before	Specifies match of files modified before this time (using the Greenwich mean time [GMT] time zone) in yyyy-mm-dd hh:mm:ss format.
time-after	Specifies match of files modified after this time (using the Greenwich mean time [GMT] time zone) in yyyy-mm-dd hh:mm:ss format.
minFileSizeInMB *minFileSizeInKB* *minFileSizeInB*	(Optional) Specifies match of content size equal to or larger than this value. The size can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).
maxFileSizeInMB *maxFileSizeInKB* *maxFileSizeInB*	(Optional) Specifies match of content size equal to or smaller than this value. The size can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).

Table B-2 <match> Subtag Attributes (continued)

Attribute	Description
prefix	(Optional.) Specifies a prefix as a match rule to filter out websites during a crawl job.
url-pattern	(Optional.) Specifies a regular expression as a match rule to filter out certain URLs.

A <match> subtag can specify multiple attributes. Attributes within a <match> tag have a Boolean AND relationship. In the following example, to satisfy this match rule, a file must have an .mpg type file extension and its size must be larger than 50 kilobytes.

```
<match extension="mpg" minFileSizeIn-KB="50" />
```

There is a Boolean OR relationship between the <match> rules themselves. A <matchRule> tag can have multiple <match> subtags, but only one of these subtags must be matched. The <matchRule> tag can be specified as a subtag of the <crawler> tag, or a subtag of the <item-group> tag. If there is a subtag in an <item-group> tag, it is shared by every <crawler> tag within that <item-group> tag.

Note The *accept* or *reject* attributes can be mistakenly used in the <crawler> tag for a crawler filter.

For example, to crawl files with the extension .mpg, simply specifying accept="\.mpg" is not correct. In this case, although specifying accept="\.mpg" is not technically incorrect, no crawling occurs. Pages with URLs that do not match the *accept* constraint are not searched. For example, if the starting URL is index.html, this HTML file is parsed and any links not containing .mpg are rejected. If the .mpg files are located in the second or lower link levels, they are not fetched because the links connecting them have been rejected.

To properly crawl for the .mpg extension, use <matchRule>. Specify <matchRule> <match extension="mpg" />. The whole site is crawled and only those files with the .mpg extension are retained.

Specifying Content Priority

A priority can be assigned to content objects to define their order of importance. The CDS software determines the order of processing from the level of priority of the content. The higher the content priority, the sooner the acquisition of content from the origin server and the sooner the content is distributed to the Service Engines.

Note Every content object acquired by running a crawl job has the same priority.

Three factors combine to determine content priority:

- Delivery Service priority—Content Distribution Priority drop-down list in the Acquisition and Distribution Properties area of the Delivery Service Definition page in the CDSM
- Item index—Content order listed in the Manifest file
- Item priority—Priority of the attributes specified in the <item> or <crawler> tag

To calculate content priority, use one of the following formulas:

- If there is a priority value for this content specified in the Manifest file *priority* attribute, use the following formula:

  Content priority = Delivery service priority * 10000 + Item priority

  In this formula, Item priority can be any integer and is unrestricted.

Tip If you want a particular content object to have the highest priority, specify a very large integer value for item priority in the content priority formula.

- If an object does not have a priority value specified in the Manifest file *priority* attribute, use the following formula:

  Content priority = Delivery service priority * 10000 + 10000 – Item index

  In this formula, Item index is the order in which content is listed in the Manifest file.

Note If there is no priority specified for any items, content is processed in the order listed in the Manifest file.

Generating a Playserver List

The CDS software supports playservers that play back the following prefetched content types on the CDS network: HTTP, HTTPS, RTSP, and RTMP (Movie Streamer, Windows Media, Flash Media Streaming).

Note Flash Media Streaming (RTMP) is a Release 2.1 feature; therefore, this feature is not available in Release 2.0.

The CDS software checks whether the requested protocol matches the list in the playserver table. If it matches, the request is delivered. If it does not match, the request is rejected.

You can generate a playserver list in the following ways:

- By configuring playserver attributes in an <item> tag
- By configuring playserver MIME-type extension names in a <playServerTable> tag

To create the playserver list directly through the Manifest file, configure playserver attributes of the playserver list in an <item> tag. If an <item> tag does not have a playserver attribute, its playserver list is generated through the <playServerTable> tag. If the <playServerTable> tag is omitted in the Manifest file, a built-in default <playServerTable> tag is used to generate the playserver list. Multiple servers are separated by commas, as shown in the following example:

```
<item src="video.mpg" playServer="wmt,http" />
```

You can also generate the playserver list that supports these streaming media types through the <playServerTable> tag. The <playServerTable> tag maps content into a playserver list based on the MIME-type extension name. If there is a <playServerTable> tag in the Manifest file, use that tag.

To generate the playserver list though the <playServerTable> tag, use MIME-type extension names to configure which playserver can play the particular prefetched content, as shown in the following example:

```
<playServerTable>
```

```
<playServer name="wmt">
         <extension name="wmv" />
         <extension name="wma" />
         <extension name="wmx" />
         <extension name="asf" />
</playServer>
<playServer name="http">
         <contentType name="application/pdf" />
         <contentType name="application/postscript" />
         <extension name="pdf" />
         <extension name="ps" />
</playServer>
</playServerTable>
```

The <playServerTable> tag is used to generate a playserver list for each content type. In the preceding example, any Portable Document Format (.pdf) or PostScript (.ps) file uses HTTP to play the content.

Customized Manifest Playserver Tables and the HTTP Playserver

In general, you do not need to specify your own playserver table or playserver in the Manifest file. A default playserver table maps appropriate file extensions or MIME-types to the proper playservers

When you use the default playserver table, the HTTP playserver is always included in the playserver list, and this allows prefetched content to be played using HTTP. If the default playserver table does not meet your needs, you can customize your playserver lists by defining your own playserver table or by specifying a *playServer* attribute in the Manifest file.

The HTTP playserver is included in the default playserver table. However, if you specify your own playserver table or *playServer* attribute in the <item> or <crawler> tags, you must add the HTTP playserver in order to play HTTP content or other content using HTTP.

Specifying Attributes for Content Serving

Certain attributes in the Manifest file can be specified to control the manner in which content is served by the Service Engines. These attributes can be specified in the <item> and <crawler> tags. These same attributes can also be specified in the <item-group> or <options> tags, so they can be shared by their <item> and <crawler> subtags. Table B-3 lists and describes these content-serving attributes.

Table B-3 Attributes for Content Serving

Attribute	Description
serveStartTime	(Optional) Designates a time in yyyy-mm-dd hh:mm:ss format at which the CDS software is allowed to start serving the content. If the serving start time is omitted, content is ready to serve once it is distributed to the Service Engine.
serveStopTime	(Optional) Designates a time in yyyy-mm-dd hh:mm:ss format at which the CDS software temporarily stops serving the content. If the serving stop time is omitted, the CDS software serves the content to the Service Engine until the content is removed by modifying the Manifest file or renaming the delivery service.

Table B-3 Attributes for Content Serving (continued)

Attribute	Description
ignoreQueryString	Playback attribute that can be used with the <options>, <item-group>, <item>, and <crawler> tags. If *ignoreQueryString* is set to true, then the CDS software ignores any string after a question mark (?) in the request URL for playback. If this attribute is omitted, then the default value is false. For example, content with the request URL url=http://web-server/foo has been prefetched. If a user requests content with the URL url=http://web-server/foo?id=xxx and the *ignoreQueryString* attribute is set to false, then the CDS software will not use prefetched content from the request URL http://web-server/foo. However, if the *ignoreQueryString* attribute is set to true, then the CDS software treats the request URL http://www-server/foo?id=xxx the same as http://www-server/foo and returns prefetched content.
wmtRequireAuth	(Optional) Determines whether users need to be authenticated before the specified content is played. When *wmtRequireAuth* is set to true, the Service Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Service Engine. If this attribute is omitted, a heuristic approach is used to determine the setting: if the specified content is acquired by using a username and password, *wmtRequireAuth* is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, *wmtRequireAuth* is set to false. **Note** If *wmtRequireAuth* is true, the Origin Server field in the CDSM Content Origin page for this delivery service needs to point to the server that can authenticate the users. When users want to play back the content, the server specified in the Origin Server field is checked for authentication.

Specifying Time Values in the Manifest File

The following attributes require that you enter a time value in the format yyyy-mm-dd hh:mm:ss (year-month-day hour:minute:second):

- *prefetch*
- *serveStartTime*
- *serveStopTime*
- *expires*
- *time-before*
- *time-after*

In the Manifest file, the time string conforms to the yyyy-mm-dd hh:mm:ss format. A time zone designation can be specified optionally at the end of a time string to indicate the particular time zone used. If a time zone designation is omitted, the GMT time zone is used. Note that automatic conversion between daylight saving time and standard time within a time zone is not supported, but a special designation for daylight saving time can be used, such as PDT for Pacific daylight saving time. In the following example, the prefetch time is September 5, 2002 at 09:09:09 Pacific daylight saving time:

```
<options timeZone="PDT" />
```

```
<item src="index.html" prefetch="2002-09-05 09:09:09 PDT" />
```

Refreshing and Removing Content

Use the *ttl* (time-to-live) and *expires* attributes of the Manifest file to monitor and control the freshness of content objects, and remove them.

The *ttl* attribute is expressed in minutes and specifies how frequently the software checks the freshness of the content at the origin server. If the *ttl* attribute is specified inside an <item> tag, it applies to that item; if it is specified inside a <crawler> tag, the attribute applies to the crawl job.

For example, if you give the *ttl* attribute a value of 10, the software checks the item or crawl job every 10 minutes. If the item has been updated, then the updated file is reacquired.

Caution Sometimes a crawl job can be very large, crawling over thousands of files. The recrawl speed is 5000 files per hour for small files. It is time-consuming to recheck so many files. We strongly recommend that you specify a large *ttl* value for such crawl jobs (for example, 1440 minutes [daily]). Otherwise, the software continues to crawl the site over and over again, blocking other acquisition tasks.

If you omit the *ttl* attribute in the Manifest file, the time-to-live is assumed to be zero and the software does not recheck that item after it is acquired. A value of 0 (zero) for *ttl* means that the content is fetched only once and is never checked again unless you click the **Fetch Manifest Now** button in the CDSM or use the **acquirer start-delivery-service** EXEC command in the Content Acquirer CLI.

The **Fetch Manifest Now** button is located in the Delivery Service Content page in the CDSM. When you click this button, the software checks to see if the Manifest file has been updated, and the updated Manifest file is downloaded and reparsed. Also, regardless of whether the Manifest file has been updated, all content in the delivery service is rechecked and the updated content is downloaded.

If you assign a negative value to the *ttl* attribute, such as –1, that item is never to be rechecked. A negative *ttl* attribute value prevents the software from checking item freshness, even if you click the **Fetch Manifest Now** button or use the **acquirer start-delivery-service** command.

Note Configuring the update interval in the CDSM GUI (**Services > Service Definition > Delivery Services > Delivery Service Content**) sets the interval for checking updates to the Manifest file itself. This setting only pertains to checking the Manifest file; it does not pertain to checking the content.

The *failRetryInterval* attribute is sometimes confused with the *ttl* attribute. The fail and retry feature acts upon failed content or failed updates. If the acquisition of a single item or of some crawled content fails, the software automatically tries to refetch these failed objects after a default interval of 5 minutes. The fail and retry interval can also be specified by using the *failRetryInterval* attribute in the Manifest file.

The difference between the *failRetryInterval* attribute and the *ttl* attribute is that the *ttl* attribute is for successfully acquired content and the *failRetryInterval* attribute is for content acquisition failures. The *ttl* attribute must be specified for the software to recheck the content freshness, whereas the *failRetryInterval* attribute does not need to be specified unless you want to change the retry interval.

The *expires* attribute specifies the time the content is to be removed from the CDS network. If you do not specify a time when you set the *expires* attribute, content is stored in the CDS network until it is explicitly removed when you modify the Manifest file. The *expires* attribute uses the format yyyy-mm-dd hh:mm:ss (year-month-day hour:minute:second). In the following example, the content expires on June 12, 2003 at 2:00 p.m.

```
expires="2003-06-12 14:00:00 PST"
```

If the *expires* attribute is specified inside an <item> tag, it applies to that item; if it is specified inside a <crawler> tag, the attribute applies to the crawl job.

You can monitor the status of content replication and freshness by enabling and then viewing the transaction log files that reside on the Service Engines. To verify whether or not a content object or file was successfully imported to or refreshed on a particular Service Engine, take these actions:

- Enable the transaction log function on the Service Engine you want to monitor.
- View the transaction log entries for the content object or filename that resides on that Service Engine.

Specifying Live Content

Only Windows Media live contents can be specified in the Manifest file. Use the <item> tag and specify the *type* attribute as wmt-live, as shown in the following example. The live stream for the wmt-live content type is url=rtsp://www.company-web-site.org/tmp/ceo-talk.

```
<CdnManifest>
<server name="wmt-server">
<host name="rtsp://www.company-web-site.org" />
</server>
<item src="/tmp/ceo-talk" type="wmt-live" >
</item>
<!--
This is a "wmt-live" streaming content type specified by the "type" attribute. The live
stream URL is
rtsp://www.company-web-site.org/tmp/ceo-talk.
-->
</CdnManifest>
```

Specifying Hybrid Ingest Content

For hybrid ingested content, the content is not prefetched into the CDS network. Instead, the content is ingested dynamically based on the user request. This type of ingest is called *dynamic ingest* or *on-demand ingest*. In order to control the play back of the on-demand content, a new type of ingest has been introduced called *hybrid ingest*. In this method, the metadata for on-demand contents can be specified in the Manifest file. However, the actual content is not acquired by the Content Acquirer.

Hybrid ingest is supported by specifying “cache” as the value for the *type* attribute inside the <item> tag.

Note This mode of ingest is supported only for single items; crawling is not supported.

Following is an example of a Manifest file for hybrid ingest content:

```
<CdnManifest>
<server name="web-server">
<host name="http://www.company-web-site.org" />
</server>
<item src="/tmp/ceo-talk.wmv" type="cache"
    serveStartTime="2007-01-12 14:00:00 PST"
    serveStopTime="2007-04-12 14:00:00 PST"
>
</item>
</CdnManifest>
```

Note For type="cache", <host> and <server> tags are not used.

Note Currently, only *serveStartTime* and *serveStopTime* are supported for type=”cache.”

Manifest Validator Utility

Because correct Manifest file syntax is so important to the proper deployment of prefetched content on your CDS network, Cisco makes available a Manifest file syntax validator. The Manifest Validator, a Java-based command-line interface that verifies the correctness of the syntax of the Manifest file you have written or modified, is built into the CDSM.

The Manifest Validator utility tests each line of the Manifest file to identify syntax errors where they exist and determine whether or not the Manifest file is valid and ready for use in importing content into your CDS network. The results of these syntax validation tests are logged into a text file at a location that you name.

Running the Manifest Validator Utility

To access the Manifest Validator, do the following:

Step 1 Choose **Services > Service Definition > Delivery Services > Tools > Manifest Validator**.

Note You must first create a new delivery service or edit an exiting delivery service before you can access the Manifest Validator.

Step 2 In the **Manifest File** field, enter the URL of the Manifest file you want to test.

Step 3 Click **Validate**.

The Manifest Validator checks the syntax of your Manifest file to make sure that source files are named for each content item in the Manifest file. It then checks the URL for each content item to verify that the content is placed correctly and then displays the output in the lower part of the page. The Manifest Validator does not determine the size of the item.

Alternatively, click **Validate** in the Delivery Service Content page. The results are displayed in a new window.

Valid Manifest File Example

The following text is an example of a valid Manifest file:

```
<CdnManifest>
<item
        src="tmp/mao's.html"
        priority="20"
        />
```

```
<server name="my-dev'box">
<host name="http://128.107.150.26"
        proto="http" />
</server>

<item
        src="tmp/lu.html"
        priority="300"
        />
<item
        src="/tmp/first_grader.html"
        />
<server name="server0">
       <host name="http://umark-u5.cisco.com:8080/" />
</server>
<item  •src="a.gif"/>
<server name="server1">
       <host name="http://unicorn-web" />
</server>
<item  •src="Media/wmtfiles/DCA%20Disk%201/Microsoft_Logos/Logos_100k.wmv" />

</CdnManifest>
```

The final lines of the Manifest Validator output indicate whether the Manifest file is valid or not. Wait until the following message is displayed, indicating that the validator has completed processing the Manifest file:

```
Total Number of Error: 0
Total Number of Warning: 0
Manifest File is CORRECT.
```

If errors are found, the error messages reported appear before the preceding message.

Invalid Manifest File Example

The following text is an example of an invalid Manifest file:

```
<CdnManifest>
<item
        src="tmp/mao's.html"
        priority="20"
        />
<server name="my-dev'box">
<host name="http://128.107.150.26"
        proto="http" />
</server>
<item
        src="tmp/lu.html"
        priority="300"
        />
<item
        src="/tmp/first_grader.html"
        />
<server name="server0">
       <host name="http://umark-u5.cisco.com:8080/" >
</server>
<item src="a.gif"/>
<server name="server1">
       <host name="http://unicorn-web" />
</server>
<item src1="Media/wmtfiles/DCA%20Disk%201/Microsoft_Logos/Logos_100k.wmv" />
</CdnManifest>
```

In the preceding example, although there are no warnings, two errors are found, and this Manifest file is syntactically incorrect, as shown in the following message:

```
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by
content model
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 9 ):Expected end of tag 'host'
  Manifest File: /state/dump/tmp.xml.1040667979990
  Total Number of Error: 2
  Total Number of Warning: 0
  Manifest File is NOT CORRECT!
```

The following full-text output is an example of the invalid Manifest file after the Manifest Validator checks the file:

```
Manifest validated: http://qiwzhang-lnx/nfs-obsidian/Unicorn/my-single-bad.xml
The manifest is downloaded as /state/dump/tmp.xml.1040667979990 for validation, this file
will be removed when validation is completed.
Start CdnManifest
Start item
     priority=20
     src=tmp/mao's.html
End item

Start server
     name=my-dev'box
Start host
     name=http://128.107.150.26
     proto=http
     uuencoded=false
End host

End server

Start item
     priority=300
     src=tmp/lu.html
End item

Start item
     src=/tmp/first_grader.html
End item

Start server
     name=server0
Start host
     name=http://umark-u5.cisco.com:8080/
     uuencoded=false
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by
content model
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 9 ):Expected end of tag 'host'
Manifest File: /state/dump/tmp.xml.1040667979990
Total Number of Error: 2
Total Number of Warning: 0
Manifest File is NOT CORRECT!
```

Understanding Manifest File Validator Output

The Manifest Validator messages appear below the Manifest File in the Manifest Validator page.

Each output file has a similar structure and syntax. It clearly identifies any errors or warning messages arising from incorrect Manifest file syntax. Manifest files are determined by the validator to be either:

- CORRECT—Contains possible syntax irregularities but is syntactically valid and ready for deployment on your CDS network
- INCORRECT—Contains syntax errors and is unsuitable for deployment on your CDS network

Syntax Errors

The Manifest Validator issues syntax errors only when it cannot identify a source file for a listed content item, either because it is not listed or because it is listed using improper syntax. Files containing syntax errors are marked INCORRECT.

Syntax errors are identified in the output with the ERROR label. In addition to the label, the line and column number containing the error are provided, as well as the Manifest file attribute for which the error was issued. An error appears in the following example:

```
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by
content model
```

In the error example:

- `/state/dump/tmp.xml.1040667979990` is the Manifest file name
- `line: 23 col: 1` is the Manifest file line and column number where the error occurs
- `No character data is allowed by content model` describes the type of Manifest file error

Syntax Warnings

The Manifest Validator issues syntax warnings for a wide variety of irregularities in the Manifest file syntax. Files containing syntax warnings may be marked CORRECT or INCORRECT, depending on whether or not syntax errors have also been issued.

Syntax warnings are identified in the output with the WARNING label. In addition to this warning label, the line number for which the warning is issued is provided, as well as the Manifest file attribute, valid options, and the default value for that attribute for which the warning was issued.

Correcting Manifest File Syntax

Once you have identified syntax warnings, errors, and messages using the output from the Manifest Validator, you can correct your Manifest file syntax and then rerun the Manifest Validator on the corrected file to verify its correctness.

It is a good idea to review every warning and error in your Manifest file. Some warnings, although they still allow the Manifest Validator to find your Manifest file syntax to be correct, can be the source of problems when you deploy the identified content to your CDS network.

Manifest File Structure and Syntax

The Cisco Manifest file provides powerful features for representing and manipulating CDS network data that can be easily edited using any simple text editor. Table B-4 provides a summary list of the Manifest file tags, their corresponding attributes and subelements, and a brief description of each tag. Table B-5 shows an example of how tags are nested in a Manifest file. The sections that follow provide a more detailed description of the Manifest file tags, the data they contain, and their attributes.

Table B-4 Manifest File Tag Summary

Tag Name	Subelements	Attributes		Description
CdnManifest	<playServerTable/> <options/> <server/> <item/> <item-group/> <crawler/>	None		Marks the beginning and end of the Manifest file content.
playServerTable	<playServer/>	None		(Optional) Sets default mappings for media types.
playServer	<contentType/> <extension/>	***name***[1]		Names the media server type on the Service Engine responsible for playing content types and files with extensions mapped to it using <contentType> tags.
contentType	None	***name***		(Optional, but must have either <contentType> or <extension> tag.) Names the MIME-type content mapped to a playserver.
extension	None	***name***		(Optional, but must have either <contentType> or <extension> tag.) Names the file extension that is mapped to a playserver.
options	<schedule/> <repeat/>	*enableCookies* *expires* *failRetryInterval* *ignoreOriginPort* *ignoreQueryString*	*prefetch* *priority* *wmtRequireAuth* *server* *sslAuthtype* *timeZone* *ttl* *type*	(Optional) Defines attributes specific to the Manifest file that can be shared.
server	<host/>	***name***		Defines *only* one host from which content is to be retrieved.
host	None	***name*** *disableBasicAuth* *noProxy* *ntlmUserDomain* *password* *port* *proto*	 *proxyServer* *sslAuthType* *user* *userDomainName* *uuencoded*	Defines a web server or live server from which content is to be retrieved and later prefetched. The hostname can be specified as: proto://user:password@hostname :port
proxyServer	None	***serverName*** *disableBasicAuth* *ntlmUserDomain* *password*	 *port* *user* *uuencoded*	Specifies proxy server information.

Table B-4 Manifest File Tag Summary (continued)

Tag Name	Subelements	Attributes		Description
item	<contains/> <schedule/> <repeat/>	***src*** *authCookie* *cdn-url* *disableBasicAuth* *enableCookies* *expires* *failRetryInterval* *host* *ignoreOriginPort* *ignoreQueryString* *noProxy* *ntlmUserDomain* *password* *playServer* *port*	 *prefetch* *priority* *proto* *proxyServer* *server* *serveStartTime* *serveStopTime* *sslAuthType* *ttl* *type* *user* *userDomainName* *uuencoded* *wmtRequireAuth*	Identifies specific content that is to be acquired from the origin server.
crawler	<matchRule/> <schedule><repeat>	***start-url*** *accept* *authCookie* *cdnPrefix* *depth* *disableBasicAuth* *enableCookies* *expires* *externalPrefixes* *externalServers* *failRetryInterval* *host* *ignoreOriginPort* *ignoreQueryString* *keepExpiredContent* *keepFolder* *keepNoCacheContent* *keepQueryUrl* *max-number* *maxTotalSizeIn-MB* *noProxy* *ntlmUserDomain*	 *password* *playServer* *port* *prefetch* *prefix* *priority* *proto* *proxyServer* *reject* *reportBrokenLinks* *serveStartTime* *serveStopTime* *server* *srcPrefix* *sslAuthType* *ttl* *type* *user* *userDomainName* *uuencoded* *wmRequireAuth*	Supports crawling of a website or FTP server.

Table B-4 Manifest File Tag Summary (continued)

Tag Name	Subelements	Attributes		Description
item-group	<item/> <crawler/> <item-group/>	*cdnPrefix* *cdn-url* *disableBasicAuth* *enableCookies* *expires* *failRetryInterval* *host* *ignoreOriginPort* *ignoreQueryString* *noProxy* *password* *playServer* *prefetch* *priority*	*proto* *proxyServer* *requireAuth* *serveStartTime* *serveStopTime* *server* *src* *srcPrefix* *sslAuthType* *ttl* *type* *user* *userDomainName* *uuencoded* *wmtRequireAuth*	Places shared attributes under one tag so that they can be shared by every <item> and <crawler> tag within that group.
matchRule	<match>	None		(Optional) Defines additional filter rules for crawler jobs.
match	None	*extension* *mime-type* *prefix* *minFileSizeIn-B* *minFileSizeIn-KB* *minFileSizeIn-MB*	*maxFileSizeIn-B* *maxFileSizeIn-KB* *maxFileSizeIn-MB* *time-after* *time-before* *url-pattern*	(Optional) Specifies the acquisition criteria of content objects before they can be acquired by the CDS network.
contains	None	*cdn-url*		(Optional) Identifies content objects that are embedded within the content item currently being described.

1. Attributes that are required for a tag are shown in ***boldface italic*** font.

Table B-5 Manifest File Nested Tag Relationships

<CdnManifest>				
	<playServerTable> <playServer>			
		<contentType /> <extension />		
			</playServerTable> </playServer>	
	<options>			
		Manifest file shared attributes	</options>	
	<server>			
		<host/>		

Table B-5 Manifest File Nested Tag Relationships (continued)

			</server>	
	<item>			
		<contains />		
			</item>	
	<crawler>			
		<matchRule/>		
			</crawler>	
	<item-group>			
		<contains />		
			</item-group>	
				</CdnManifest>

CdnManifest

The<CdnManifest> </CdnManifest> tag set is required and marks the beginning and end of the Manifest file content. At a minimum, each <CdnManifest> tag set must contain at least one item, or content object, that is fetched and stored.

Attributes

None

Subelements

The <CdnManifest> tag set can contain the following subelements:

- playServerTable

  The <CdnManifest> tag set can contain only one playServerTable subelement.
- options

  The <CdnManifest> tag set can contain only one options subelement.
- server
- item
- item-group
- crawler

Example

```
<CdnManifest>
    <server name="origin-server">
                <host name="www.name.com" proto="http" port="80" />
    </server>
    <item cdn-url= "logo.jpg" server="originserver"  src= "images/img.jpg" type="prepos"
              playServer="http" ttl="300"/>
</CdnManifest>
```

playServerTable

The <playServerTable> </playServerTable> tag set is optional and provides a means for you to set default mappings for a variety of media types. Mappings can be set for both MIME-type content (the preferred mapping) and file extensions. Playserver tables allow you to override default mappings on the Service Engine for content types from a particular origin server. Playservers can be any one of the following streaming servers: WMT, HTTP, QTSS, or FMS. If no <playServerTable> tag is configured in the Manifest file, a default <playServerTable> tag is used.

Note Flash Media Streaming (FMS) is a Release 2.1 feature; therefore, this feature is not available in Release 2.0.

Using the Manifest file, you can map groups of single items as well as individual content objects to an installed playserver. The following are content item and Manifest file playserver mappings:

- Content item URL

  Playserver mappings appear immediately after the origin server name in place of the default <CdnManifest> tag.

- Manifest file as an attribute of the <item> or <item-group> tag

  Playserver mappings placed at this location are identified using the *playServer* attribute and only apply to the named item or group of items.

- Manifest file as a playserver table

  Mappings are grouped within the <playServerTable> and <playServer> tags and are applied to content served from the origin server as directed by the Manifest file.

- System-level

  Playserver mappings are configured during CDS software startup.

The <playServerTable> tags are enclosed within the <CdnManifest> tags and name at least one of four playservers, such as RealServer, to which certain MIME-types and file extensions are mapped.

Attributes

None

Subelements

The <playServerTable> element must contain at least one <playServer> tag.

playServer

The <playServer> </playServer> tag set is required for the <playServerTable> tag and names the media server type on the Service Engine that is responsible for playing the content types and files with extensions mapped to it using the <contentType> tags. The <playServer> tag is enclosed within <playServerTable> tags.

Note Do not confuse the <playServer> tag with the *playserver* attribute in an <item> or <item-group> tag. An <item> or <item-group> tag specifies a server type to be used for an individual content object or group of related content objects. Although both playserver settings accomplish the same task, <item> tag-level playserver settings take precedence over the content type and file extension mappings specified by the <playServer> tags in the <playServerTable> tag.

Attributes

The <playServer> tag name is required. Each <playServer> tag names the type of server to which content is mapped using the *name* attribute. The Service Engines support the following types of playservers:

- http: HTTP web server
- qtss: Apple QuickTime Streaming Server
- wmt: Microsoft Windows Media Technologies
- fms: Flash Media Streaming Server

Note Flash Media Streaming is a Release 2.1 feature; therefore, this feature is not available in Release 2.0.

Subelements

At least one of the following subelements must be present in a <playServer> tag set.

- <contentType />
- <extension />

contentType

The <contentType /> tag is optional but either a <contentType /> or an <extension /> subelement must be present in a <playServer> tag set. The <contentType /> tag names MIME-type content that is to be mapped to a playserver. The <contentType /> tag must be enclosed within a <playServer> tag set. When both <contentType /> and <extension /> tags are present in a <PlayServerTable> tag for a particular media type, the <contentType /> mapping takes precedence.

Attributes

Each <contentType /> tag names a media content type that is to be mapped to the playserver using the *name* attribute. The *name* attribute is required.

Subelements

None

extension

The <extension /> tag is optional but either a <contentType /> or an <extension /> subelement must be present in a <playServer> tag set. The <extension /> tag names the file extension that is being mapped to a playserver.

The <extension /> tag follows the <playServer> tag. When both <contentType /> and <extension /> tags are present in the <playServer> tag for a particular media type, the <contentType /> mapping takes precedence.

Attributes

The *name* attribute is required and provides the file extension for a mapped content type. When files with the named extension are requested, the mapped playserver is used to serve them.

Subelements

None

Example

```
<CdnManifest>
<playServerTable>
<playServer name="wmt">
         <extension name="asf" />
</playServer>
<playServer name="http">
         <contentType name="application/pdf" />
         <contentType name="application/postscript" />
         <extension name="pdf" />
         <extension name="ps" />
</playServer>
</playServerTable>
<server name="test.origin.com/">
         <host name="http://tst.orgn.com" proto="http" />
</server>
<item
    src="pic1.mpg"
/>
</CdnManifest>
```

options

The <options/> tag is optional and used to define attributes specific to the Manifest file. Shared attributes can be inherited by <item> and <crawler> tags in the Manifest file. For example, *timeZone* is an attribute specific to the Manifest file that is used to set the time zone for all time-related values. Attributes such as *ttl* can exist as <options/> tags, and their values can be shared by all <item> and <crawler> tags within the Manifest file.

The <options/> tag set is enclosed within the <CdnManifest> tag set and specifies at least one global setting. No more than one <options> tag is allowed per Manifest file.

If parameters are defined within the Manifest file <options/>, <item-group>, or <item> tags, the order of precedence from lowest to highest is <options/>, <item-group>, and <item>.

Attributes

The *timeZone* attribute specifies the time zone for time values of attributes such as *expires* and *prefetch*.

The following list of attributes can be shared by <item> and <crawler> tags. See the "item" section on page B-28 for descriptions of the following attributes:

- *enableCookies*
- *expires*
- *failRetryInterval*
- *ignoreOriginPort*
- *ignoreQueryString*
- *prefetch*

- *priority*
- *wmtRequireAuth*
- *server*
- *sslAuthType*
- *ttl*
- *type*

Subelements

<schedule><repeat>

(See the "item" section on page B-28 for descriptions of these subelements.)

server

The <server> and <host> tag fields configure the origin content source server. The <host> tag field inside the <server> tag field configures the content source host. Having multiple <host> tag fields in one <server> tag field is not supported.

Each <item> or <item-group> tag can have a *server* attribute that refers to this <server> tag field. The <server> </server> tag set is required and defines only one host from which content is to be retrieved. The <server> tags are contained within <CdnManifest> tags and contain one <host> tag that identifies the host from which content is retrieved.

Attributes

The *name* attribute is required and can be any name as long as it matches the *server* attribute values in the <item> or <crawler> tags.

Subelements

The <server> tag set can only contain one <host/> subelement.

host

The <host/> tag is required and defines a web server or live server from which content is to be retrieved and later prefetched. Only one host can be defined within a single <server> tag set. The <host/> tag must be enclosed within <server> tags.

Attributes

- *disableBasicAuth*

  The *disableBasicAuth* attribute is optional; if specified, basic authentication is disabled.

- *name*

  The *name* attribute is required and identifies the domain name or IP address of the host, unless the *proto* attribute field is empty. If the *proto* attribute field is empty, the *name* attribute must be a fully qualified URL, including scheme and domain name or IP address. It can also include subdirectories, such as http://www.abc.com/media.

  The *name* attribute can also contain the UNC path to an SMB server; for example, \\SMBserver\directory\.

- *noProxy*

The *noProxy* attribute is optional. If set to true, no proxy is used for the origin server. The default is false.

- *ntlmUserDomain*

  The *ntlmUserDomain* attribute is optional and specifies the user domain name for NTLM authentication.

- *password*

  The *password* attribute is optional and identifies the password for the user account that is required to access the host server.

- *port*

  The *port* attribute is optional and identifies the TCP port through which traffic to and from the host passes. The port used depends on the protocol used. The default port for HTTP is 80. The *port* attribute is only required for a nonstandard port assignment. The port attribute can also be specified in the *name* attribute, such as name="http://www.cisco.com:8080/."

- *proto*

  The *proto* attribute is optional and identifies the communication protocol that is used to fetch content from the host. Supported protocols are HTTP, HTTPS, MMS-over-HTTP, or FTP. The default *proto* attribute is HTTP. The *proto* attribute can be empty if the *name* attribute is a fully qualified domain name (FQDN).

- *proxyServer*

  The *proxyServer* attribute is optional and specifies which proxy server to use if there are multiple <proxyServer> tags in the Manifest file. If no proxy server is specified, the server in the closest <proxyServer> tag is used.

- *sslAuthType*

  The *sslAuthType* attribute is optional and has two possible values for the type of SSL certificate verification:

  - *strong*—Strong authentication. If any errors occur during certificate verification by the acquirer module, content from that site is not acquired. The default *sslAuthType* attribute setting is strong.
  - *weak*—Weak authentication. If certain errors occur during certificate verification by the acquirer module, content from that site continues to be acquired. These errors are as follows:

    Unable to decode issuer's public key

    Certificate has expired

    Self-signed certificate

    Self-signed certificate in certificate chain

    Unable to get local issuer certificate

    Subject issuer mismatch

    Authority and issuer serial number mismatch

    The Content Acquirer is not marked as trusted

    Unable to verify the first certificate

    Certificate is not yet valid

    Certificate has invalid purpose

- *user*

  The *user* attribute is optional and identifies the secure login used for host access.

- *userDomainName*

  See the "item" section on page B-28 for a description of this attribute.

- *uuencoded*

  The *uuencoded* attribute is optional. If set to true, the password is not encoded. The *uuencoded* attribute default setting is false.

Subelements

None

proxyServer

The <proxyServer> tag specifies proxy server information. The <proxyServer> tag must be located at the top level of the Manifest file, directly under the <CdnManifest> tag; it cannot be used as a subtag of any other tags, as shown in this example:

```
<CdnManifest>
<proxyServer>
...
</CdnManifest>
</proxyServer>
```

Attributes

- *disableBasicAuth*

  The *disableBasicAuth* attribute is optional; if specified, basic authentication is disabled.

- *ntlmUserDomain*

  The *ntlmUserDomain* attribute is optional and specifies the user domain name for NTLM authentication.

- *password*

  The *password* attribute is optional and identifies the password for the user account that is required to access the proxy server.

- *port*

  The *port* attribute is optional and specifies the proxy port.

- *serverName*

  The *serverName* attribute is required and identifies the domain name or IP address of the proxy server.

- *user*

  The *user* attribute is optional and identifies the secure login used for proxy authentication.

- *uuencoded*

  The *uuencoded* attribute is optional and designates whether the password is to be encoded.

Subelements

None

item

The <item> </item> tag set identifies the specific content that is to be acquired. The <item> tag names a single piece of content or a content object on the origin server, such as a graphic, MPEG video, or RealAudio sound file. Content items can be listed individually or grouped using the <item-group> tag.

The <item> tag must be enclosed within the <CdnManifest> tag set and can also be enclosed within <item-group> tags.

Attributes

- *src*

  The *src* attribute is required and identifies the URL from which to fetch the content. The URL can be a full URL or a relative URL. A full URL has the following format:

  proto://username:password@/domain-name:port/file-path/file-name

  Protocols supported in the *src* attribute are HTTP, HTTPS, FTP, and SMB. For SMB, the URL must be written in UNC format (\\SMBserver\directory\file).

  If a relative path is used, the <server> and <host> tags are required to specify origin server information, as shown in this example:

```
<item src="http://user:password@www.cisco.com/HR/index.html" />
<server name="ftp-server" >
   <host name="ftp://ftp-server"  user="johw" password="wwww" />
</host>
<item src="data/video.asf" />
```

  Note A URL containing a question mark (?) is not supported. A Manifest file parsing error will occur if you specify a URL that contains a question mark.

  Note A URL containing a pound sign (#) will be modified. All characters that follow a pound sign will be discarded, including the pound sign itself.

- *host*

  The *host* attribute specifies the hostname if the source URL of the *src* attribute is a relative URL.

- *server*

  The *server* attribute is optional and refers to the server name in the <server> tag. If the *server* attribute is omitted, the server listed in the closest <server> tag is used. If there is no <server> tag close to this item, the Manifest file server is used.

- *cdn-url*

  The *cdn-url* attribute is optional and is used when content needs to be acquired from one URL (the content acquisition URL) and published using another URL (the publishing URL). The *cdn-url* attribute is the relative CDS network URL that end users use to access this content. If no *cdn-url* attribute is specified, then the *src* attribute is used as the relative CDS network URL.

  In the following sample Manifest file, the content item being acquired contains the file path /RemAdmin/InternalReview/firstpage.htm. By specifying a new file path (RemAdmin/Production/firstpage.htm) using the *cdn-url* attribute, the publishing URL disguises the fact that the content originated from an “Internal Review.”

```
<CdnManifest>
```

```
<server name="ultra-server">
    <host name="http://ultra-server" />
</server>
<item src="RemAdmin/InternalReview/firstpage.htm"
cdn-url="RemAdmin/Production/firstpage.htm" />
</CdnManifest>
```

In the preceding example, *src is* the content acquisition URL and *cdn-url is* the publishing URL.

Note The content item file path (RemAdmin/InternalReview/firstpage.htm) is controlled by the Manifest file. The *cdn-url* attribute associates a file path with the content item in the Manifest file. The Manifest file allows the file path for the *cdn-url* attribute to be specified independently of the file path from which the content items are to be acquired from the origin server (*src* attribute), allowing the publishing URL to differ from the content acquisition URL.

If the content requires playback authentication or is live content, the origin server from which the content is acquired has to be contacted. Therefore, two URLs must exist for the same content item, and the URL specified in the *cdn-url* attribute must exist on the origin server at all times.

For example, if the content item "RemAdmin/Production/firstpage.htm" in the preceding example requires playback authentication, this content must exist on the "ultra-server" origin server. Otherwise, prefetched content playback will fail.

In general, you should not use the *cdn-url*, *cdnPrefix*, or *srcPrefix* attributes if playback authentication is required or if the content is live.

If you use FTP to acquire content and the content type is not specified in the Manifest file and the *cdn-url* attribute is specified to alter your publishing URL, the *cdn-url* attribute must have the correct file path extension. Otherwise, the incorrect content type will be generated and you cannot play the content.

The following example correctly shows the publishing URL with the same file path extension (.jpg) as the origin server URL.

```
<item src="ftp://ftp-server.abc.com/pictures/pic.jpg"  cdn-url="pic.jpg" />
```

The following example is incorrectly written, because it does not specify the file path extension (.jpg) in the *cdn-url* attribute.

```
<item src="ftp://ftp-server.abc.com/pictures/pic.jpg"  cdn-url="pic" />
```

- *type*

  The *type* attribute is optional and defines whether content is to be prefetched or live on the CDS network. The three type attributes are *prepos*, *cache*, and *wmt-live*. The *wmt-live* type attribute is used to deliver live content. The *cache* type corresponds to hybrid ingest method. If this field is left blank, the default type is *prepos*.

  Note For type="cache", <host> and <server> tags are not used.

  Note Currently, only *serveStartTime* and *serveStopTime* are supported for the type="cache" attribute.

- *playServer*

The *playServer* attribute is optional and names the server used to play back the content. Valid playservers are wmt (Windows Media Technologies), qtss (QuickTime Streaming Server), fms (Flash Media Streaming), and http (Web Engine). The value in this field is either one playserver or multiple playservers separated by commas. If a value for this attribute is not specified, the <PlayServerTable> tag in the Manifest file is used to generate the playserver list for this content. If the Manifest file does not have the <PlayServerTable> tag specified, it uses the default <PlayServerTable> tag.

- *prefetch*

  The *prefetch* attribute is optional and specifies a time (in yyyy-mm-dd hh:mm:ss [year-month-day hour:minute:second] format) for the first content acquisition or re-check after the Manifest file is parsed. The time zone for the time can be specified in the <options> tag. Note that the autoconversion between daylight saving time and standard time within a time zone is not supported, but a special designation for daylight saving time can be used, such as PDT for Pacific daylight saving time. In the following example, the prefetch time is September 5, 2002 at 09:09:09 Pacific daylight saving time.

  ```
  <options timeZone="PDT" />
  <item src="index.html" prefetch="2002-09-05 09:09:09 PDT" />
  ```

  This attribute is used when you want to specify a future time for the acquirer to begin fetching content from the origin server. When a future time is specified, the acquirer will not acquire content before this time; however, it will check content freshness during its scheduled *ttl* interval. If a *prefetch* time is omitted, the content is acquired immediately.

  After the Manifest file is parsed, if any items or crawl tasks have changed or new ones have been added and if the *prefetch* attribute specifies a future time, the acquirer checks and fetches the content or re-crawls the crawl jobs at the time specified by the *prefetch* attribute.

- *expires*

  The *expires* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the content is to be removed from the CDS network. Additionally, you can specify the GMT time zone. If a time value is omitted, content is stored until it is removed when you modify the relevant Manifest file code.

- *ttl*

  The *ttl* attribute is optional and designates a time interval, in minutes, for revalidation of the content. If a time value is omitted, the content is fetched only once and its freshness is never checked again. Usually the *ttl* attribute is a positive value; however, you can also assign a negative value to the *ttl* attribute. The following table describes *ttl* attribute value ranges.

***ttl* Attribute Value**	**Action**
ttl > 0	Content is rechecked every *ttl* minute. Content is also rechecked if the Manifest file is reparsed and the content specification in the Manifest file has changed or if you click the **Refetch** button.
ttl = 0	Content is fetched only once and never checked again. Content is only rechecked if the Manifest file is reparsed and the content specification in the Manifest file has changed or if you click the **Refetch** button.
ttl < 0	Content is fetched only once and never checked again. Content will *not* be rechecked if the Manifest file is reparsed or if you click the **Refetch** button.

- *serveStartTime*

The *serveStartTime* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the CDS software is allowed to start serving the content. If the time to serve is omitted, content is ready to serve once it is distributed to the Service Engine or other edge device.

- *serveStopTime*

  The *serveStopTime* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the CDS software temporarily stops serving the content. If the time to stop serving is omitted, the CDS software serves the content until it is removed when you modify the relevant Manifest file code.

- *priority*

  The *priority* attribute is optional and can be any integer value to specify the content processing priority. If a priority value is omitted, its index order within the Manifest file is used to set the priority.

- *wmtRequireAuth*

  The *wmtRequireAuth* attribute is optional and determines whether users need to be authenticated before the specified content is played. When true, the Service Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Service Engine. If this attribute is omitted, a heuristic approach is used to determine the value: if the specified content is acquired by using a username and password, *wmtRequireAuth* is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, *wmtRequireAuth* is set to false.

  Note If *wmtRequireAuth* is true, the Origin Server field in the Content Origin page for this delivery service needs to point to the server that can authenticate users. When users want to play back the content, the server specified in the Origin Server field is checked for authentication.

- *failRetryInterval*

  The *failRetyryInterval* attribute specifies the retry interval, in minutes, when content acquisition fails. For example, failRetryInterval="10" means the CDS software retries content acquisition every 10 minutes after acquisition has failed. If the retry universal value is not specified, the default value is 5 minutes. (The minimum *failRetryInterval* value is accepted.) If a value of less than 5 minutes is specified, that value is converted to 5 minutes.

  The behavior differs between failed content acquisition of a single item and failed content acquisition of a crawl item.

  – For single item failure:

    ```
    if ( ttl != 0, ttl < retryInterval)
    ```

    The item is rechecked in accordance with the *ttl* attribute. Otherwise, the item is rechecked at the interval specified in the *failRetryInterval* attribute.

  – For crawl item failure:

    ```
    if ( ttl != 0 and ttl < retryInterval )
    always re-crawl
    ```

    If some items are not acquired (excluding 300 and 400 series status error codes), only failed items are rechecked as specified in the *failRetryInterval* attribute.

    When the *ttl* attribute interval occurs, all pages are recrawled.

  For example, if *ttl* = 10, and *failRetryInterval* = 4, the following actions occurs:

Number of Minutes	Action
0	Crawl
4	Recheck failed
8	Recheck failed
10	Recrawl
14	Recheck
18	Recheck
20	Recrawl

- *ignoreQueryString*

  The *ignoreQueryString* attribute is a playback attribute that can be used with the <options>, <item-group>, <item>, and <crawler> tags. If the value is set to true, then CDS software ignores any string after a question mark (?) in the request URL for playback. If this attribute is omitted, then the default value is false.

  For example, content with the request URL url=http://web-server/foo has been prefetched. If a user requests content with the URL url=http://web-server/foo?id=xxx and the *ignoreQueryString* attribute value is false, then CDS software does not use the prefetched content from the request URL http://web-server/foo.

  However, if the *ignoreQueryString* attribute is set to true, then the CDS software treats the request URL http://www-server/foo?id=xxx the same as http://www-server/foo and returns with prefetched content.

- *ignoreOriginPort*

  The *ignoreOriginPort* attribute allows playback of prefetched content from a port other than the standard port. If the *ignoreOriginPort* attribute is set to true, content can be played back without regard to the port specified in the request URL. The default for this attribute is false.

  This attribute is not intended to be used for content that is routed using a Service Router. It is intended to work only for explicit proxy routing. A typical usage scenario for the *ignoreOriginPort* attribute might be as follows:

  - The origin web server is not using port 80; it is using a nonstandard port number in the URL.
  - Users are using explicit proxy routing, where the original URL containing the non-standard port number is used for playback from the Service Engine.

  Prefetched content cannot be played back using a nonstandard port; prefetched content is served only on ports that are standard for the protocol. If the incoming URL contains a port number other than the protocol's standard port, you must set the *ignoreOriginPort* attribute to true for playback to succeed.

- *userDomainName*

  The *userDomainName* attribute is used in two instances: for NTLM authentication and for the SMB file import feature. If the origin server is using NTLM authentication, you must use this attribute to specify the user domain name for NTLM authentication. If a shared folder is protected and the user account is part of a domain, you must use this attribute to specify the domain name of the configured shared folder.

Note Both *userDomainName* and *ntlmUserDomain* cannot coexist in the Manifest file; only one attribute can be used at a time.

- *enableCookies*

  The *enableCookies* attribute enables cookie support for the item. When this attribute is set to true, the Content Acquirer, after sending a request for an item to the origin server, parses the server response for cookie name/value pairs. If the server response contains a cookie that is valid and has not expired, the Content Acquirer stores the cookie in main memory.

  The Content Acquirer then returns the valid cookie to the server the next time the Content Acquirer sends a request for the item.

  A cookie is rejected if it contains any of the following rejection criteria, as found in RFC 2965:

  – The value for the Path is not a prefix of the request URI.
  For example, if the request is www.abc.com/aaa/bbb/ccc.html and the Path of the cookie returned is /aaa/ccc, then it is not valid because /aaa/ccc is not a prefix of /aaa/bbb/ccc [URL].

  – The value for the Domain contains no embedded dots or does not start with a dot.

  – The value for the request host is not a domain-match of the Domain.

  – The request host is a FQDN (not an IP address) and has the form HD, where D is the value of the Domain and H is a string that contains one or more dots.

  – The Path is not a prefix match of the request URL.

  Note The Content Acquirer does not use persistent memory to store cookies. If the Service Engine is restarted, all cookie information is lost.

  The *enableCookies* attribute can be used with the <item>, <crawler>, <item-group>, and <options> tags.

- *authCookie*

  The *authCookie* attribute enables the processing and sending of authentication cookies for the item. To enable this feature, the *authCookie* attribute must be set to true for the particular item that passes the user credentials and for which the server sends back the authentication cookies.

  The *authCookie* attribute can be used with the <item> and <crawler> tags. For example:

  <item src=http://abc.com/auth.cgi?id=10000 authCookie="true"/>

The following attributes described under the <host> tag attributes can also be specified by the <item> tag.

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*
- *port*
- *proto*
- *proxyServer*
- *sslAuthType*

- *user*
- *uuencoded*

Subelements

- <contains />
- <schedule/> <repeat/>

  The <schedule/> <repeat/> subelement and its attributes specify a time for a recrawl or an item refetch to begin. You can have multiple <repeat> subelements under the <schedule> subelement. The attributes *time, start,* and *end* specify the day of the month or day of the week and the duration of the specified repeat. The *time* attribute is required, whereas *start* and *end* are optional attributes.

  Note The <schedule> element takes precedence over the *ttl* attribute.

  The *time* attribute uses either of the following formats:

  time=“dom:hh:mm” or

  time=“dow:hh:mm”

  In these formats, dom is the day of the month (0–30), dow is the day of the week (Sun, Mon, Tue, Wed, Thu, Fri, Sat, or *), hh is the clock hour (0–23 or *), and mm is the minute (0–59).

  For example:

```
<schedule>
    <repeat time="*:*:0" /> <!-- repeat every hour on the hour -->
    <repeat time="*:13:0" /><!-- repeat at 1300 every day -->
    <repeat time="Sun:2:30" /> <!-- repeat on Sundays at 2:30 -->
    <repeat time="4:2:30" /> <!-- repeat at 2:30 on the fourth day of the month -->
    <repeat time="Mon:*:30" /> <!-- On Monday, repeat every hour at 30 minutes past
        the hour -->
</schedule>
```

  The *start* and *end* attributes use the following format:

  start=“yyyy-mm-dd hh:mm:ss”

  end=“yyyy-mm-dd hh:mm:ss”

  For example:

```
<CdnManifest>
<item>
    <schedule>
        <repeat time="Sun:02:30" />
        <repeat time="*:*:34" start="2003-09-11 11:11:11 PST" end="2004-09-11 11:11:21
PST"/>
        <repeat time="21:02:35" start="2003-09-11 11:11:11 PST" end="2004-09-11
11:11:21 PST"/>
        <repeat time="21:02:35" end="2004-09-11 11:11:21 PST"/>
    </schedule>
</item>
.
.
.
```

Example

```
<item
    src="index.html"
    server="cisco.com"
```

```
    ttl="3000"
/>
```

crawler

The <crawler> </crawler> tag set supports crawling a website or an FTP server.

Attributes

- *start-url*

  The *start-url* attribute is required. It defines the URL at which to start the process of crawling the website or FTP server. It is identical to the *src* attribute used in the <item> tag. (See the " •src" section on page B-28 under the item section.)

- *host*

  The *host* attribute specifies the host name if the starting URL specified in the *start-url* attribute is a relative URL.

- *depth*

  The *depth* attribute is optional and defines the link depth to which a website is to be crawled or directory depth to which an FTP server is to be crawled. If the depth is not specified, the default is 20. The following are the general depth values:

  0 = Acquire only the starting URL
  1, 2, 3, ... = Acquire the starting URL and its referred files
  –1 = Infinite or no depth restriction

  Depth is defined as the level of a website or the directory level of an FTP server, where 0 is the starting URL.

- *prefix*

  The *prefix* attribute is optional and combines the hostname from the <server> tag with the value of the *prefix* attribute to create a full prefix. Only content with URLs that match the full prefix is acquired, as shown in this example:

  ```
  <server name="xx"> <host name="www.cisco.com" proto="https" port=433 /> </server>
  ```

  and with the following <crawler> tag:

  ```
  prefix="marketing/eng/"
  ```

  The full prefix is "https://www.cisco.com:433/marketing/eng/." Only URLs that match this prefix are crawled.

  If a prefix is omitted, the crawler checks the default full prefix, which is the hostname portion of the URL from the server. In the example, the default full prefix is "https://www.cisco.com:433."

- *accept*

  The *accept* attribute is optional and uses a regular expression to define acceptable URLs to crawl in addition to matching the prefix. For example, accept="stock" means that only URLs that meet two conditions are searched: the URL matches the prefix and contains the string "stock." (See the "Writing Common Regular Expressions" section on page B-5 for more information on using regular expressions.)

  Note the following two key differences between the *accept* attribute and the *prefix* attribute:

  – The *prefix* attribute uses an exact string match, while the *accept* attribute uses a regular expression.

  - The *prefix* attribute applies to a URL including all of its links or subdirectories. However, the *accept* attribute allows the URL and its links and subdirectories to be evaluated separately.

- *reject*

  The *reject* attribute is optional and uses a regular expression to reject a URL if it matches the reject regular expression. The reject regular expression is checked after checking for a prefix URL match. If a URL does not match the prefix, it is immediately rejected. If a URL matches the prefix and the reject parameters, it is rejected by the particular reject constraint. (See the "Writing Common Regular Expressions" section on page B-5 for more information on using regular expressions.)

  Note the following two key differences between the *reject* attribute and the *prefix* attribute:

  - The *prefix* attribute uses an exact string match, while the *reject* attribute uses a regular expression.
  - The *prefix* attribute applies to a URL including all of its links or subdirectories. However, the *reject* attribute allows the URL and its links and subdirectories to be evaluated separately.

- *max-number*

  The *max-number* attribute is optional and specifies the maximum number of crawler job objects that can be acquired.

- *maxTotalSizeInB/KB/MB*

  The *maxTotalSizeInB/KB/MB* attribute is optional and specifies the maximum total content size in bytes, kilobytes, or megabytes that this crawler job can acquire. The size attribute can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).

  This attribute replaces the *max-size-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *srcPrefix*

  The *srcPrefix* attribute is optional and must be used in conjunction with the *cdnPrefix* attribute to form a relative CDS network URL. If a *srcPrefix* attribute is not specified, or if the prefix of the relative source URL does not match the *srcPrefix* attribute, then the relative CDS network URL is the *cdnPrefix* value combined with the relative source URL. For example, if these content objects have the same source URL prefix "acme/pubs/docs/online/Design/" and you want to replace this prefix with a simple "online/," then specify srcPrefix="acme/pubs/docs/online/Design/" and cdnPrefix="online/."

- *cdnPrefix*

  The *cdnPrefix* attribute is optional and must be used in conjunction with the *srcPrefix* attribute.

- *wmtRequireAuth*

  The *wmtRequireAuth* attribute is optional and determines whether users need to be authenticated before the specified content is played. When true, the Service Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Service Engine. If this attribute is omitted, a process of discovery approach is used to determine the value: if the specified content is acquired by using a username and password, *wmtRequireAuth* is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, *wmtRequireAuth* is set to false.

Note If *wmtRequireAuth* is set to true, the Origin Server field in the Content Origin page for this delivery service needs to point to the server that can authenticate the users. When users want to play back the content, the server specified in the Origin Server field is checked for authentication.

- *externalPrefixes*

  The *externalPrefixes* attribute is optional and specifies additional prefixes for crawl jobs to crawl multiple protocols or multiple websites. Prefixes are separated with a bar (|).

- *externalServers*

  The *externalServers* attribute is optional and can be used for multiple host crawling jobs where each host has a different user account. This attribute can be used to refer to the <host> tag with the proper authentication information.

- *keepExpiredContent*

  The *keepExpiredContent* attribute can be used to acquire content during an HTTP or HTTPS crawl that is expired. When this attribute is set to true, expired content will be fetched. When this attribute is set to false, expired content is discarded. If this attribute is not specified, the default is false.

- *keepFolder*

  The *keepFolder* attribute is used to fetch folders (a folder is indicated when the request URL ends with a forward slash "/"). If this attribute is set to false, folder URLs are not acquired.

- *keepNoCacheContent*

  The *keepNoCacheContent* attribute can be used to acquire content during an HTTP or HTTPS crawl that would normally not be cached. When this attribute is set to true, the acquirer will fetch the content even though the content contains an HTTP cache control header indicating that the content is not to be cached. If this attribute is not specified, the default is false.

- *keepQueryUrl*

  The *keepQueryUrl* attribute can be used to fetch URLs that contain "?" in the URL string. If this attribute is set to true, URLs with "?" will be fetched during HTML parsing for a crawl job if the URL meets the other crawling criteria set forth in the Manifest file.

  This attribute is useful when you want to acquire content from a database, for example, where multiple files are differentiated in the portion of the URL string after the "?". When this attribute is not set, the portion of the URL after the "?" is discarded. If multiple URLs are found where the portion of the URL string in front of the "?" is the same, these URLs appear as duplicates, and only the last "duplicate" URL found is fetched.

- *reportBrokenLinks*

  The *reportBrokenLinks* attribute is used to report links on an HTML web page that cannot be fetched. If this attribute is set to true, all broken links encountered during a website crawl will be reported as errors. This attribute only applies to a website crawl, not to an index crawl. The default is false and broken links are not reported as errors.

The following attributes described under the <host> tag attributes can also be specified by the <crawler> tag:

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*
- *port*
- *proto*
- *proxyServer*
- *sslAuthType*

- *user*
- *uuencoded*

The following attributes described under the <item> tag attributes can also be specified by the <crawler> tag.

- *authCookie*
- *enableCookies*
- *expires*
- *failRetryInterval*
- *ignoreOriginPort*
- *ignoreQueryString*
- *playServer*
- *prefetch*
- *priority*
- *serveStartTime*
- *serveStopTime*
- *server*
- *ttl*
- *type*
- *userDomainName*

Subelements

- <matchRule></matchRule>
- <schedule><repeat>

  (See the "item" section on page B-28 for descriptions of the <schedule><repeat> subelements.)

Example

```
<server name="cisco">
    <host name="http://www.cisco.com/jobs/" />
</server>
<crawler
    server="cisco"
    start-url="eng/index.html"
    depth="10"
    prefix="eng/"
    reject="\.pl"
    maxTotalSizeIn-MB="200"
/>
```

item-group

The <item-group> </item-group> tag set is used to place shared attributes under one tag so that they can be shared by every <item> and <crawler> tag within that group. When attributes are shared, it means that attributes can be defined at either the <item-group> tag level for group-wide control or on a per

<item> or per <crawler> tag basis. For example, if every <item> tag is using the same *server* and *ttl* attributes, you can create an <item-group> tag on top of these <item> tags and place the *server* and *ttl* attributes in the <item-group> tag.

Using shared attributes makes any Manifest file with many <item> tags more efficient by consolidating the <item> tags with shared attributes. If the same attribute value exists in both the <item-group> and <item> tags, the value in the <item> tag takes precedence over that value in the <item-group> tag.

The <item-group> tag must be enclosed within the <CdnManifest> tag set and contain one or more <item> or <crawler> tags.

Attributes

If an attribute value is present only at the <item-group> tag level, then it is inherited by its inner element in the <item> tag. If an attribute value is present in a crawler job, its attributes, whether inherited or owned, are propagated to the content fetched by the crawler job.

The following attributes can be shared across many <item> and <crawler> tags and are candidates for the <item-group> level tag. See the "item" section on page B-28 for detailed descriptions of the following attributes:

- *cdn-url*
- *enableCookies*
- *expires*
- *failRetryInterval*
- *host*
- *ignoreOriginPort*
- *ignoreQueryString*
- *playServer*
- *prefetch*
- *wmtRequireAuth*
- *serveStartTime*
- *serveStopTime*
- *server*
- *src*
- *priority*
- *ttl*
- *type*
- *userDomainName*

The following attributes described under the <host> tag attributes can also be specified by the <item-group> tag.

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*
- *port*

- *proto*
- *proxyServer*
- *sslAuthType*
- *user*
- *uuencoded*

Additionally, the following two attributes can be placed within the <item-group> tag. See the "crawler" section on page B-35 for a detailed description of the following two attributes:

- *srcPrefix*
- *cdnPrefix*

These two attributes convert the prefix of the *src-url* (content acquisition URL) to the *cdn-url* (publishing URL) for multiple content objects. These content objects are either implicitly specified by multiple <item> tags or acquired through a crawler job.

These two attributes can also be specified in the <crawler> tag. If you explicitly specify the *srcPrefix* attribute and *cdnPrefix* attribute for an individual <crawler> job, the <crawler> tag-level specification takes precedence over the <item-group> tag-level settings. If you do not specify these attributes for an individual <crawler> job, the <item-group> tag-level specification is inherited by the <crawler> job.

The *srcPrefix* and *cdnPrefix* attributes generate the relative CDS network URL using the following rules:

- If the *cdn-url* attribute is present in the <item> tag, the relative CDS network URL contains both the *cdnPrefix* attribute plus the *cdn-url* attribute. For example, if *cdnPrefix*="eng/spec" and *cdn-url*="e/f.html," the relative path in the URL is "eng/spec/e/f.html."
- If the *srcPrefix* attribute is not present in the <item> tag, the relative CDS network URL is the *cdnPrefix* attribute plus the relative source URL.
- If the prefix of the relative source URL does not match the *srcPrefix* attribute, the relative CDS network URL is the *cdnPrefix* attribute plus the relative source URL.
- To generate a relative CDS network URL, remove the matched prefix from the relative source URL and replace it with the *cdnPrefix* attribute.

The relative CDS network URL of <item> in the following example is "acme/default.htm."

```
<item-group cdnPrefix="acme/" >
    <item src="design/index.html" cdn-url="default.html" />
</item-group>
```

In the following example, content objects with the *srcPrefix* attribute, such as "design/plan/," have the relative CDS network URL as "acme/" plus relative source URLs stripped of "design/plan/." Other content objects with a prefix attribute that does not match "design/plan/" have "acme/" plus their original relative source URL.

```
<crawler
    start-url="design/plan/index.html"
    depth="-1"
    srcPrefix="design/plan/"
    cdnPrefix="acme/" />
```

Subelements

- <crawler></crawler>
- <item-group/>
- <item></item>
- <schedule><repeat>

(See the "item" section on page B-28 for descriptions of the <schedule><repeat> subelements.)

Example

```
<!--grouped content items-->
<item-group server="origin-web-server" type="prepos" ttl="300" cdnPrefix="unicorn/" >
         <item cdn-url="newHQpresentation.rm" src="newHQpresentation.rm" />
         <item cdn-url="animatedlogo.mpg" src="animlogo.mpg" />
         <item cdn-url="companytheme.mp3" src="cotheme.mp3" />
         <item cdn-url="newHQlayout.avi" src="newHQ.mov" />
</item-group>
```

matchRule

The <matchRule> </matchRule> tag set is optional and defines additional filter rules for crawler jobs. It affects only <crawler> tasks and is not used by single <item> tags. The crawler parameters defined in the <crawler></crawler> tag set determine primarily the scope of a crawl search. If a content object does not meet the criteria specified by the crawler parameter, neither it nor its children are searched.

The <matchRule> tag, however, determines only whether or not the content objects should be acquired regardless of the scope of the search. If a web page matches the crawler parameters without the <matchRule> feature, its children are searched even though its content objects are not acquired.

In the following crawler job example that uses the <matchRule> tag, the entire website is searched, but only files with the .jpg file extension larger than 50 kilobytes are acquired.

```
<crawler  start-url="index.html"  depth="-1" >
    <matchRule>
        <match minFileSizeIn-KB="50" extension="jpg" />
    </matchRule>
</crawler>
```

The <matchRule> element can be nested within an <item-group> tag to define group-wide filter rules for <crawler> tags contained in the group. It can also be a subelement of a particular <crawler> job. The <crawler> tag-level setting overrides the <item-group> tag-level setting when both tags are present.

If you define criteria locally for individual <crawler> jobs, any existing group-level criterion is entirely discarded for that <crawler> job. If your <item-group> tag match rule is set to A and your <crawler> tag specifies another match rule set to B, only B is to be used for the <crawler> tag rather than a combination of A and B. You can define at most one <matchRule> tag per <item-group> tag and at most one <matchRule> tag per <crawler> tag.

Attributes

None

Subelements

At least one <match> tag

match

The <match> </match> tag is optional and specifies the acquisition criteria of content objects before they can be acquired by CDS software. Every attribute within a single <match> tag has a Boolean AND relationship (to form a logical conjunction) with the other attributes.

You can specify multiple <match> tags within the <matchRule> tag. The <match> tags have a Boolean OR relationship (to form a logical inclusion) with other <match> tags. You must specify at least one <match> tag per <matchRule> tag.

Attributes

- *mime-type*

  The *mime-type* attribute specifies MIME-types.

- *extension*

  The *extension* attribute specifies file extensions.

- *time-before*

  The *time-before* attribute can provide both an absolute time (modified before yyyy-mm-dd hh:mm:ss) or a relative time (modified within ddd:hh:ss), relative to the present time, to download content. Time parameters should be expressed in GMT time zones. (For GMT offsets, see the "Manifest File Time Zone Tables" section on page B-50.)

- *time-after*

  The *time-after* attribute can provide both an absolute time (modified after yyyy-mm-dd hh:mm:ss) or a relative time (modified within ddd:hh:ss), relative to the present time, to download content. Time parameters should be expressed in GMT time zones. (For GMT offsets, see the "Manifest File Time Zone Tables" section on page B-50.)

  Note Relative time is calculated based on current time. We recommend that you synchronize the the server clock and the Service Engine clock so that relative time calculations are accurate.

- *minFileSizeInB/KB/MB*

  The *minFileSizeInB/KB/MB* attribute specifies that the acquired content size must be larger than this number of bytes, kilobytes, or megabytes. The size attribute can be expressed in bytes (B), kilobytes (KB), or megabytes (MB).

  The *minFileSizeInB/KB/MB* attribute replaces the *size-min-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *maxFileSizeInB/KB/MB*

  The *maxFileSizeInB/KB/MB* attribute specifies that the acquired content size must be smaller than this number of bytes, kilobytes, or megabytes. This attribute can be expressed in bytes (B), kilobytes (KB), or megabytes (MB).

  The *maxFileSizeInB/KB/MB* attribute replaces the *size-max-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *prefix*

  The *prefix* attribute is optional and specifies a prefix as a match rule to filter out websites during a crawl job.

- *url-pattern*

  The *url-pattern* attribute is optional and specifies a regular expression as a match rule to filter out certain URLs.

Subelements

None

Examples

```
<! - - crawling item group -- >
<item-group server="origin-server" type="prepos">
       <matchRule>
              <match time-before="2000-05-05 12:0:0"/>
       </matchRule>
          <crawler start-url="eng/index.html" depth="-1"/>
            <crawler start-url="hr/index.html" depth="3">
                       <matchRule>
                              <match minFileSizeIn-KB="1" extension="xxx"/>
                       </matchRule>
            </crawler>
</item-group>
```

To download content that was created or modified within the last 90 days, use the relative time format, as shown in the following example:

```
<match time-after="90:00:00"/>
```

To download content that was not modified within the last 2 weeks, use the relative time format, as shown in the following example:

```
<match time-before="14:00:00"/>
```

To download content that has been modified after January 30, 2003, 10:30 p.m., use the absolute time format, as shown in the following example:

```
<match time-after="2003-01-30 10:30:00"/>
```

contains

The <contains /> tag is optional and identifies content objects that are embedded within the content item currently being described. For example, the components of a Synchronized Multimedia Integration Language (SMIL) file request for an item using <contains /> links are only accepted after CDS software determines that dependent content objects are present in the Service Engine.

The <contains /> tag must be enclosed within the <item> </item> tag.

The <contains /> tag is used to include embedded files for some video files, such as .asf or .rp. The CDS software does not serve this item unless every contained item is present.

Attributes

The *cdn-url* attribute is required and is the relative CDS network URL of one of the embedded contents.

Subelements

None

Example

```
<item src="house/img08.jpb" cdn-url="img08.jpg" />
<item src="house/img09.jpb" cdn-url="img09.jpg" />
<item cdn-url="house.rp"src="house/house.rp">
     <contains cdn-url="img08.jpg"/>
     <contains cdn-url="img09.jpg"/>
</item>
```

XML Schema

In the case of the Manifest file, an XML schema defines the custom markup language of the Manifest file and the appearance of a given set of XML documents. The XML schema specifies which tags or elements you can use in your documents, the attributes those tags can contain, and their arrangement.

Manifest XML Schema

An XSD is a library that provides an application programming interface (API) for manipulating the components of an XML schema.

The following XML code is the Manifest XML schema (CdnManfest.xsd):

```
<?xml version="1.0"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
<xs:include schemaLocation="PlayServerTable.xsd"/>
<xs:element name="CdnManifest">
  <xs:complexType>
    <xs:sequence>
        <xs:element ref="playServerTable" minOccurs="0" maxOccurs="1"/>
        <xs:element ref="options" minOccurs="0" maxOccurs="1"/>
        <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element ref="proxyServer" maxOccurs="unbounded"/>
        <xs:element ref="server" maxOccurs="unbounded"/>
            <xs:element ref="item-group" maxOccurs="unbounded"/>
            <xs:element ref="item" maxOccurs="unbounded"/>
            <xs:element ref="crawler" maxOccurs="unbounded"/>
        </xs:choice>
        <xs:choice minOccurs="0" maxOccurs="unbounded">
            <xs:element ref="item-group" maxOccurs="unbounded"/>
            <xs:element ref="item" maxOccurs="unbounded"/>
            <xs:element ref="crawler" maxOccurs="unbounded"/>
        </xs:choice>
   </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="options">
    <xs:complexType>
         <xs:sequence>
             <xs:element ref="schedule" minOccurs="0" maxOccurs="1"/>
         </xs:sequence>
        <xs:attribute name="timeZone" type="xs:string" use="optional"/>
        <xs:attribute name="wmtRequireAuth" type="xs:boolean" use="optional"/>
        <xs:attribute name="ignoreQueryString" type="xs:boolean" use="optional"/>
        <xs:attribute name="ignoreOriginPort" type="xs:boolean" use="optional"/>
        <xs:attribute name="enableCookies" type="xs:boolean" use="optional" />
        <xs:attribute name="ttl" type="xs:int" use="optional"/>
        <xs:attribute name="failRetryInterval" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="prefetch" type="xs:string" use="optional"/>
        <xs:attribute name="type" use="optional" >
            <xs:simpleType>
                <xs:restriction base="xs:string">
                    <xs:enumeration value="prepos"/>
                    <xs:enumeration value="wmt-live"/>
                    <xs:enumeration value="cache"/>
                 </xs:restriction>
            </xs:simpleType>
        </xs:attribute>
        <xs:attribute name="manifest-id" type="xs:string" use="optional"/>
```

```
        <xs:attribute name="clearlog" type="xs:boolean" use="optional" />
        <xs:attribute name="rd" type="xs:string" use="optional"/>
        <xs:attribute name="prepos-tag" type="xs:string" use="optional"/>
        <xs:attribute name="live-tag" type="xs:string" use="optional"/>
    </xs:complexType>
</xs:element>

<xs:element name="server">
    <xs:complexType>
        <xs:sequence>
            <xs:element ref="host" minOccurs="1" maxOccurs="1"/>
        </xs:sequence>
        <xs:attribute name="name" type="xs:string" use="required"/>
    </xs:complexType>
</xs:element>

<xs:attributeGroup name = "hostAttr">
    <xs:attribute name="proxyServer" type="xs:string" use="optional"/>
    <xs:attribute name="noProxy" type="xs:boolean" use="optional"/>
    <xs:attribute name="proto" use="optional">
        <xs:simpleType>
              <xs:restriction base="xs:string">
                  <xs:enumeration value="http"/>
                  <xs:enumeration value="https"/>
                  <xs:enumeration value="ftp"/>
                  <xs:enumeration value="rtsp"/>
                  <xs:enumeration value="file"/>
              </xs:restriction>
        </xs:simpleType>
    </xs:attribute>
    <xs:attribute name="port" type="xs:unsignedShort" use="optional"/>
    <xs:attribute name="user" type="xs:string" use="optional"/>
    <xs:attribute name="password" type="xs:string" use="optional"/>
    <xs:attribute name="ntlmUserDomain" type="xs:string" use="optional"/>
    <xs:attribute name="userDomainName" type="xs:string" use="optional"/>
    <xs:attribute name="disableBasicAuth" type="xs:boolean" use="optional"/>
    <xs:attribute name="uuencoded" type="xs:boolean" use="optional"/>
    <xs:attribute name="sslAuthType" use="optional">
        <xs:simpleType>
              <xs:restriction base="xs:string">
                  <xs:enumeration value="weak"/>
          <xs:enumeration value="strong"/>
          </xs:restriction>
        </xs:simpleType>
   </xs:attribute>
</xs:attributeGroup>

<xs:element name="host">
    <xs:complexType>
          <xs:attribute name="name" type="xs:string" use="required"/>
          <xs:attributeGroup ref="hostAttr"/>
 </xs:complexType>
</xs:element>

<xs:element name="proxyServer">
    <xs:complexType>
        <xs:attribute name="serverName" type="xs:string" use="required"/>
        <xs:attribute name="host" type="xs:string" use="optional"/>
    <xs:attribute name="port" type="xs:unsignedShort" use="optional"/>
        <xs:attribute name="user" type="xs:string" use="optional"/>
        <xs:attribute name="password" type="xs:string" use="optional"/>
    <xs:attribute name="ntlmUserDomain" type="xs:string" use="optional"/>
        <xs:attribute name="uuencoded" type="xs:string" use="optional" />
    <xs:attribute name="disableBasicAuth" type="xs:boolean" use="optional" />
```

```
    </xs:complexType>
</xs:element>

<xs:attributeGroup name = "contentAttr">
    <xs:attribute name="server" type="xs:string" use="optional"/>
    <xs:attribute name="playServer" type="xs:string" use="optional"/>
    <xs:attribute name="type" use="optional">
        <xs:simpleType>
            <xs:restriction base="xs:string">
                <xs:enumeration value="prepos"/>
                <xs:enumeration value="wmt-live"/>
                <xs:enumeration value="cache"/>
            </xs:restriction>
        </xs:simpleType>
    </xs:attribute>

    <xs:attribute name="enableCookies" type="xs:boolean" use="optional" />
    <xs:attribute name="wmtRequireAuth" type="xs:boolean" use="optional"/>
    <xs:attribute name="ignoreQueryString" type="xs:boolean" use="optional"/>
    <xs:attribute name="ignoreOriginPort" type="xs:boolean" use="optional"/>
    <xs:attribute name="ttl" type="xs:int" use="optional"/>
    <xs:attribute name="failRetryInterval" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="priority" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="prefetch" type="xs:string" use="optional"/>
    <xs:attribute name="expires" type="xs:string" use="optional"/>
    <xs:attribute name="serve" type="xs:string" use="optional"/>
    <xs:attribute name="serveStartTime" type="xs:string" use="optional"/>
    <xs:attribute name="serveStopTime" type="xs:string" use="optional"/>
    <xs:attribute name="_isDefault" type="xs:boolean" use="optional"/>
</xs:attributeGroup>

<xs:attributeGroup name = "prefixAttr">
     <xs:attribute name="cdnPrefix" type="xs:string" use="optional"/>
     <xs:attribute name="srcPrefix" type="xs:string" use="optional"/>
</xs:attributeGroup>

<xs:element name="item-group">
    <xs:complexType>
        <xs:sequence>
            <xs:element ref="matchRule" minOccurs="0" maxOccurs="1"/>
            <xs:element ref="schedule" minOccurs="0" maxOccurs="1"/>
                <xs:choice minOccurs="1" maxOccurs="unbounded">
                <xs:element ref="item-group" maxOccurs="unbounded"/>
                <xs:element ref="item" maxOccurs="unbounded"/>
                <xs:element ref="crawler" maxOccurs="unbounded"/>
            </xs:choice>
        </xs:sequence>
        <xs:attributeGroup ref="contentAttr"/>
        <xs:attributeGroup ref="prefixAttr"/>
        <xs:attribute name="host" type="xs:string" use="optional"/>
        <xs:attributeGroup ref="hostAttr"/>
    </xs:complexType>
</xs:element>

<xs:element name="item">
    <xs:complexType>
        <xs:sequence>
            <xs:element ref="contains" minOccurs="0" maxOccurs="unbounded"/>
            <xs:element ref="schedule" minOccurs="0" maxOccurs="1"/>
            </xs:sequence>
    <xs:attribute name="src" type="xs:string" use="required"/>
        <xs:attribute name="cdn-url" type="xs:string" use="optional"/>
        <xs:attributeGroup ref="contentAttr"/>
    <xs:attribute name="host" type="xs:string" use="optional"/>
```

```
    <xs:attribute name="authCookie" type="xs:boolean" use="optional" />
        <xs:attributeGroup ref="hostAttr"/>
    </xs:complexType>
</xs:element>

<xs:element name="crawler">
    <xs:complexType>
        <xs:all>
            <xs:element ref="matchRule" minOccurs="0" maxOccurs="1"/>
            <xs:element ref="schedule" minOccurs="0" maxOccurs="1"/>
        </xs:all>
        <xs:attribute name="start-url" type="xs:string" use="required"/>
        <xs:attribute name="externalPrefixes" type="xs:string" use="optional"/>
        <xs:attribute name="externalServers" type="xs:string" use="optional"/>
        <xs:attribute name="keepFolder" type="xs:boolean" use="optional" />
        <xs:attribute name="keepQueryUrl" type="xs:boolean" use="optional" />
        <xs:attribute name="keepNoCacheContent" type="xs:boolean" use="optional" />
        <xs:attribute name="keepExpiredContent" type="xs:boolean" use="optional" />
        <xs:attribute name="depth" type="xs:short" use="optional"/>
        <xs:attribute name="prefix" type="xs:string" use="optional"/>
        <xs:attribute name="accept" type="xs:string" use="optional"/>
        <xs:attribute name="reject" type="xs:string" use="optional"/>
        <xs:attribute name="max-number" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="reportBrokenLinks" type="xs:boolean" use="optional" />
        <xs:attribute name="maxTotalSizeIn-B" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="maxTotalSizeIn-KB" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="maxTotalSizeIn-MB" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="maxTotalSizeInB" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="maxTotalSizeInKB" type="xs:unsignedInt" use="optional"/>
        <xs:attribute name="maxTotalSizeInMB" type="xs:unsignedInt" use="optional"/>
        <xs:attributeGroup ref="contentAttr"/>
        <xs:attributeGroup ref="prefixAttr"/>
    <xs:attribute name="host" type="xs:string" use="optional"/>
    <xs:attribute name="authCookie" type="xs:boolean" use="optional" />
        <xs:attributeGroup ref="hostAttr"/>
    </xs:complexType>
</xs:element>

<xs:element name="contains">
    <xs:complexType>
        <xs:attribute name="cdn-url" type="xs:string" use="required"/>
    </xs:complexType>
</xs:element>

<xs:element name="matchRule">
    <xs:complexType>
        <xs:sequence>
        <xs:element ref="match" minOccurs="1" maxOccurs="unbounded"/>
        </xs:sequence>
    </xs:complexType>
</xs:element>

<xs:element name="match">
    <xs:complexType>
        <xs:attribute name="mime-type" type="xs:string" use="optional"/>
        <xs:attribute name="time-before" type="xs:string" use="optional"/>
        <xs:attribute name="time-after" type="xs:string" use="optional"/>
        <xs:attribute name="minFileSizeIn-B" type="xs:int" use="optional"/>
        <xs:attribute name="maxFileSizeIn-B" type="xs:int" use="optional"/>
        <xs:attribute name="minFileSizeIn-KB" type="xs:int" use="optional"/>
        <xs:attribute name="maxFileSizeIn-KB" type="xs:int" use="optional"/>
        <xs:attribute name="minFileSizeIn-MB" type="xs:int" use="optional"/>
        <xs:attribute name="maxFileSizeIn-MB" type="xs:int" use="optional"/>
        <xs:attribute name="minFileSizeInB" type="xs:int" use="optional"/>
```

```
        <xs:attribute name="maxFileSizeInB" type="xs:int" use="optional"/>
        <xs:attribute name="minFileSizeInKB" type="xs:int" use="optional"/>
        <xs:attribute name="maxFileSizeInKB" type="xs:int" use="optional"/>
        <xs:attribute name="minFileSizeInMB" type="xs:int" use="optional"/>
        <xs:attribute name="maxFileSizeInMB" type="xs:int" use="optional"/>
        <xs:attribute name="extension" type="xs:string" use="optional"/>
        <xs:attribute name="prefix" type="xs:string" use="optional"/>
        <xs:attribute name="url-pattern" type="xs:string" use="optional"/>
    </xs:complexType>
 </xs:element>


 <xs:element name="schedule">
    <xs:complexType>
        <xs:sequence>
        <xs:element ref="repeat" minOccurs="1" maxOccurs="unbounded"/>
        </xs:sequence>
    </xs:complexType>

 </xs:element>

 <xs:element name="repeat">
    <xs:complexType>
        <xs:attribute name="time" type="xs:string" use="required"/>
        <xs:attribute name="start" type="xs:string" use="optional"/>
        <xs:attribute name="end" type="xs:string" use="optional"/>
    </xs:complexType>
 </xs:element>
</xs:schema>
```

PlayServerTable XML Schema

The following XML code defines the PlayServerTable schema (playServerTable.xsd) for the CdnManfiest.xsd:

```
<?xml version="1.0"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
<xs:element name="playServerTable">
     <xs:complexType>
         <xs:sequence>
             <xs:element ref="playServer" minOccurs="1" maxOccurs="unbounded"/>
         </xs:sequence>
     </xs:complexType>
 </xs:element>
<xs:element name="playServer">
     <xs:complexType>
         <xs:choice minOccurs="1" maxOccurs="unbounded">
             <xs:element ref="contentType"/>
             <xs:element ref="extension"/>
         </xs:choice>
         <xs:attribute name="name" use="required">
             <xs:simpleType>
                 <xs:restriction base="xs:string">
                      <xs:enumeration value="wmt"/>
                      <xs:enumeration value="http"/>
                      <xs:enumeration value="qtss"/>
                 </xs:restriction>
             </xs:simpleType>
```

```
        </xs:attribute>
    </xs:complexType>
 </xs:element>
 <xs:element name="contentType">
    <xs:complexType>
        <xs:attribute name="name" type="xs:string" use="required"/>
    </xs:complexType>
 </xs:element>
 <xs:element name="extension">
    <xs:complexType>
        <xs:attribute name="name" type="xs:string" use="required"/>
    </xs:complexType>
 </xs:element>
</xs:schema>
```

Default PlayServerTable Schema

The following XML code defines the default PlayServerTable:

```
<?xml version="1.0"?>

<playServerTable xmlns:xsi = "http://www.w3.org/2001/XMLSchema-instance"
     xsi:noNamespaceSchemaLocation = "PlayServerTable.xsd">

   <!-- playServer http and https can always play all prepositon
      contents unless users use customized <playServerTable>
      or "playServer" attribute in the manifest file
      -->

  <playServer name="qtss">
    <contentType name="video/quicktime" />
    <extension name="mov" />
    <extension name="qt" />
    <extension name="mp4" />
    <extension name="3gp" />
    <extension name="3g2" />

    <!-- extension avi could also go here -->
  </playServer>

  <playServer name="wmt">
    <!-- MIME types taken from
    http://msdn.microsoft.com/workshop/imedia/windowsmedia/server/mime.asp
    -->
    <contentType name="video/x-ms-asf" />
    <contentType name="audio/x-ms-wma" />
    <contentType name="video/x-ms-wmv" />
    <contentType name="video/x-ms-wm"  />
    <contentType name="application/x-ms-wmz" />
    <contentType name="application/x-ms-wmd" />

    <extension name="wma" /> <!-- audio content -->
    <extension name="wmv" /> <!-- audio/video content -->
    <extension name="asf" /> <!-- audio/video content (legacy) -->
    <extension name="wm"  /> <!-- reserved for future use -->

   </playServer>

</playServerTable>
```

Manifest File Time Zone Tables

To convert to local time, you must know the time difference between Greenwich mean time (GMT) and local time for both standard time and summer time (daylight saving time). Table B-6 through Table B-21 list the time zones supported by the Manifest file. The format for writing the time zone is:

<*zonename*>:[+|-:]hh:mm per line

In this format, <*zonename*> is the name of the time zone or standard time zone abbreviation (see Table B-6) without spaces before or after the colon (":"), and "[+|-:]hh:mm" is the GMT offset in hours and minutes. The GMT offset default is "+."

Table B-6 Standard Time Zones and GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
ACT:+:09:30	Etc/GMT+7:-:07:00	HST:-:10:00
ADT:-:03:00	Etc/GMT+8:-:08:00	IET:-:05:00
AET:+:10:00	Etc/GMT+9:-:09:00	IST:+:05:30
AGT:-:03:00	Etc/GMT-0:00:00	JST:+:09:00
ART:+:02:00	Etc/GMT-10:+:10:00	MDT:-:06:00
AST:-:09:00	Etc/GMT-11:+:11:00	MET:+:01:00
BET:-:03:00	Etc/GMT-12:+:12:00	MIT:-:11:00
BST:+:06:00	Etc/GMT-13:+:13:00	MST7MDT:-:07:00
CAT:+:02:00	Etc/GMT-14:+:14:00	MST:-:07:00
CDT:-:05:00	Etc/GMT-1:+:01:00	NET:+:04:00
CET:+:01:00	Etc/GMT-2:+:02:00	NST:+:12:00
CNT:-:03:30	Etc/GMT-3:+:03:00	NZ-CHAT:+:12:45
CST6CDT:-:06:00	Etc/GMT-4:+:04:00	NZ:+:12:00
CST:-:06:00	Etc/GMT-5:+:05:00	Navajo:-:07:00
CTT:+:08:00	Etc/GMT-6:+:06:00	PDT:-:07:00
EAT:+:03:00	Etc/GMT-7:+:07:00	PLT:+:05:00
ECT:+:01:00	Etc/GMT-8:+:08:00	PNT:-:07:00
EDT:-:04:00	Etc/GMT-9:+:09:00	PRC:+:08:00
EET:+:02:00	Etc/GMT0:00:00	PRT:-:04:00
EST5EDT:-:05:00	Etc/GMT:00:00	PST8PDT:-:08:00
EST:-:05:00	Etc/Greenwich:00:00	PST:-:08:00
Etc/GMT+0:00:00	Etc/UCT:00:00	ROK:+:09:00
Etc/GMT+10:-:10:00	Etc/UTC:00:00	SST:+:11:00
Etc/GMT+11:-:11:00	Etc/Universal:00:00	UCT:00:00
Etc/GMT+12:-:12:00	Etc/Zulu:00:00	UTC:00:00
Etc/GMT+1:-:01:00	GB-Eire:00:00	Universal:00:00
Etc/GMT+2:-:02:00	GB:00:00	VST:+:07:00
Etc/GMT+3:-:03:00	GMT0:00:00	W-SU:+:03:00

Table B-6 Standard Time Zones and GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Etc/GMT+4:-:04:00	GMT:00:00	WET:00:00
Etc/GMT+5:-:05:00	Greenwich:00:00	Zulu:00:00
Etc/GMT+6:-:06:00	HDT:-:09:00	—

Table B-7 Africa GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Africa/Abidjan:00:00	Africa/Djibouti:+:03:00	Africa/Maputo:+:02:00
Africa/Accra:00:00	Africa/Douala:+:01:00	Africa/Maseru:+:02:00
Africa/Addis_Ababa:+:03:00	Africa/El_Aaiun:00:00	Africa/Mbabane:+:02:00
Africa/Algiers:+:01:00	Africa/Freetown:00:00	Africa/Mogadishu:+:03:00
Africa/Asmera:+:03:00	Africa/Gaborone:+:02:00	Africa/Monrovia:00:00
Africa/Bamako:00:00	Africa/Harare:+:02:00	Africa/Nairobi:+:03:00
Africa/Bangui:+:01:00	Africa/Johannesburg:+:02:00	Africa/Ndjamena:+:01:00
Africa/Banjul:00:00	Africa/Kampala:+:03:00	Africa/Niamey:+:01:00
Africa/Bissau:00:00	Africa/Khartoum:+:03:00	Africa/Nouakchott:00:00
Africa/Blantyre:+:02:00	Africa/Kigali:+:02:00	Africa/Ouagadougou:00:00
Africa/Brazzaville:+:01:00	Africa/Kinshasa:+:01:00	Africa/Porto-Novo:+:01:00
Africa/Bujumbura:+:02:00	Africa/Lagos:+:01:00	Africa/Sao_Tome:00:00
Africa/Cairo:+:02:00	Africa/Libreville:+:01:00	Africa/Timbuktu:00:00
Africa/Casablanca:00:00	Africa/Lome:00:00	Africa/Tripoli:+:02:00
Africa/Ceuta:+:01:00	Africa/Luanda:+:01:00	Africa/Tunis:+:01:00
Africa/Conakry:00:00	Africa/Lubumbashi:+:02:00	Africa/Windhoek:+:01:00
Africa/Dakar:00:00	Africa/Lusaka:+:02:00	—
Africa/Dar_es_Salaam:+:03:00	Africa/Malabo:+:01:00	—

Table B-8 America GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
America/Adak:-:10:00	America/Grenada:-:04:00	America/Noronha:-:02:00
America/Anchorage:-:09:00	America/Guadeloupe:-:04:00	America/North_Dak/Ctr:-:06:00
America/Anguilla:-:04:00	America/Guatemala:-:06:00	America/Panama:-:05:00
America/Antigua:-:04:00	America/Guayaquil:-:05:00	America/Pangnirtung:-:05:00
America/Araguaina:-:03:00	America/Guyana:-:04:00	America/Paramaribo:-:03:00
America/Aruba:-:04:00	America/Halifax:-:04:00	America/Phoenix:-:07:00
America/Asuncion:-:04:00	America/Havana:-:05:00	America/Port-au-Prince:-:05:00
America/Atka:-:10:00	America/Hermosillo:-:07:00	America/Port_of_Spain:-:04:00

Table B-8 America GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
America/Barbados:-:04:00	America/Ind/Indian:-:05:00	America/Porto_Acre:-:05:00
America/Belem:-:03:00	America/Ind/Knox:-:05:00	America/Porto_Velho:-:04:00
America/Belize:-:06:00	America/Ind/Marengo:-:05:00	America/Puerto_Rico:-:04:00
America/Boa_Vista:-:04:00	America/Ind/Vevay:-:05:00	America/Rainy_River:-:06:00
America/Bogota:-:05:00	America/Indianapolis:-:05:00	America/Rankin_Inlet:-:06:00
America/Bogota:-:05:00	America/Inuvik:-:07:00	America/Recife:-:03:00
America/Buenos_Aires:-:03:00	America/Iqaluit:-:05:00	America/Regina:-:06:00
America/Cambridge_Bay:-:07:0	America/Jamaica:-:05:00	America/Rio_Branco:-:05:00
America/Cancun:-:06:00	America/Jujuy:-:03:00	America/Rosario:-:03:00
America/Caracas:-:04:00	America/Juneau:-:09:00	America/Santiago:-:04:00
America/Catamarca:-:03:00	America/Ken/Louisville:-:05:00	America/Santo_Domingo:-:04:0
America/Cayenne:-:03:00	America/Ken/Monticello:-:05:0	America/Sao_Paulo:-:03:00
America/Cayman:-:05:00	America/Knox_IN:-:05:00	America/Scoresbysund:-:01:00
America/Chicago:-:06:00	America/La_Paz:-:04:00	America/Shiprock:-:07:00
America/Chihuahua:-:07:00	America/Lima:-:05:00	America/St_Johns:-:03:30
America/Cordoba:-:03:00	America/Los_Angeles:-:08:00	America/St_Lucia:-:04:00
America/Costa_Rica:-:06:00	America/Louisville:-:05:00	America/St_Thomas:-:04:00
America/Cuiaba:-:04:00	America/Maceio:-:03:00	America/St_Vincent:-:04:00
America/Curacao:-:04:00	America/Managua:-:06:00	America/Swift_Current:-:06:00
America/Danmarkshavn:00:00	America/Manaus:-:04:00	America/Tegucigalpa:-:06:00
America/Dawson:-:08:00	America/Martinique:-:04:00	America/Thule:-:04:00
America/Dawson_Creek:-:07:00	America/Mazatlan:-:07:00	America/Thunder_Bay:-:05:00
America/Denver:-:07:00	America/Mendoza:-:03:00	America/Tijuana:-:08:00
America/Detroit:-:05:00	America/Menominee:-:06:00	America/Tortola:-:04:00
America/Dominica:-:04:00	America/Merida:-:06:00	America/Vancouver:-:08:00
America/Edmonton:-:07:00	America/Mexico_City:-:06:00	America/St_Lucia:-:04:00
America/Eirunepe:-:05:00	America/Miquelon:-:03:00	America/Virgin:-:04:00
America/El_Salvador:-:06:00	America/Monterrey:-:06:00	America/Whitehorse:-:08:00
America/Ensenada:-:08:00	America/Montevideo:-:03:00	America/Winnipeg:-:06:00
America/Fort_Wayne:-:05:00	America/Montreal:-:05:00	America/Yakutat:-:09:00
America/Fortaleza:-:03:00	America/Montserrat:-:04:00	America/Yellowknife:-:07:00
America/Glace_Bay:-:04:00	America/Nassau:-:05:00	America/Virgin:-:04:00
America/Godthab:-:03:00	America/New_York:-:05:00	America/Whitehorse:-:08:00
America/Goose_Bay:-:04:00	America/Nipigon:-:05:00	America/Winnipeg:-:06:00
America/Grand_Turk:-:05:00	America/Nome:-:09:00	America/Tortola:-:04:00

Table B-9 Antarctica/Arctic GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Antarctica/Casey:+:08:00	Antarctica/McMurdo:+:12:00	Antarctica/Vostok:+:06:00
Antarctica/Davis:+:07:00	Antarctica/Palmer:-:04:00	Arctic/Longyearbyen:+:01:00
Antarctica/DtDUrville:+:10:00	Antarctica/South_Pole:+:12:00	—
Antarctica/Mawson:+:06:00	Antarctica/Syowa:+:03:00	—

Table B-10 Asia GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Asia/Aden:+:03:00	Asia/Hong_Kong:+:08:00	Asia/Riyadh87:+:03:07
Asia/Almaty:+:06:00	Asia/Hovd:+:07:00	Asia/Riyadh88:+:03:07
Asia/Amman:+:02:00	Asia/Irkutsk:+:08:00	Asia/Riyadh89:+:03:07
Asia/Anadyr:+:12:00	Asia/Istanbul:+:02:00	Asia/Riyadh:+:03:00
Asia/Aqtau:+:04:00	Asia/Jakarta:+:07:00	Asia/Saigon:+:07:00
Asia/Aqtobe:+:05:00	Asia/Jayapura:+:09:00	Asia/Sakhalin:+:10:00
Asia/Ashgabat:+:05:00	Asia/Jerusalem:+:02:00	Asia/Samarkand:+:05:00
Asia/Ashkhabad:+:05:00	Asia/Kabul:+:04:30	Asia/Seoul:+:09:00
Asia/Baghdad:+:03:00	Asia/Kamchatka:+:12:00	Asia/Shanghai:+:08:00
Asia/Bahrain:+:03:00	Asia/Karachi:+:05:00	Asia/Singapore:+:08:00
Asia/Baku:+:04:00	Asia/Kashgar:+:08:00	Asia/Taipei:+:08:00
Asia/Bangkok:+:07:00	Asia/Katmandu:+:05:45	Asia/Tashkent:+:05:00
Asia/Beirut:+:02:00	Asia/Krasnoyarsk:+:07:00	Asia/Tbilisi:+:04:00
Asia/Bishkek:+:05:00	Asia/Kuala_Lumpur:+:08:00	Asia/Tehran:+:03:30
Asia/Brunei:+:08:00	Asia/Kuching:+:08:00	Asia/Tel_Aviv:+:02:00
Asia/Calcutta:+:05:30	Asia/Kuwait:+:03:00	Asia/Thimbu:+:06:00
Asia/Choibalsan:+:09:00	Asia/Macao:+:08:00	Asia/Thimphu:+:06:00
Asia/Chongqing:+:08:00	Asia/Magadan:+:11:00	Asia/Tokyo:+:09:00
Asia/Chungking:+:08:00	Asia/Manila:+:08:00	Asia/Ujung_Pandang:+:08:00
Asia/Colombo:+:06:00	Asia/Muscat:+:04:00	Asia/Ulaanbaatar:+:08:00
Asia/Dacca:+:06:00	Asia/Nicosia:+:02:00	Asia/Ulan_Bator:+:08:00
Asia/Damascus:+:02:00	Asia/Novosibirsk:+:06:00	Asia/Urumqi:+:08:00
Asia/Dhaka:+:06:00	Asia/Omsk:+:06:00	Asia/Vientiane:+:07:00
Asia/Dili:+:09:00	Asia/Phnom_Penh:+:07:00	Asia/Vladivostok:+:10:00
Asia/Dubai:+:04:00	Asia/Pontianak:+:07:00	Asia/Yakutsk:+:09:00
Asia/Dushanbe:+:05:00	Asia/Pyongyang:+:09:00	Asia/Yekaterinburg:+:05:00
Asia/Gaza:+:02:00	Asia/Qatar:+:03:00	Asia/Yerevan:+:04:00
Asia/Harbin:+:08:00	Asia/Rangoon:+:06:30	—

Table B-11 Atlantic GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Atlantic/Azores:-:01:00	Atlantic/Faeroe:00:00	Atlantic/South_Georgia:-:02:00
Atlantic/Bermuda:-:04:00	Atlantic/Jan_Mayen:+:01:00	Atlantic/St_Helena:00:00
Atlantic/Canary:00:00	Atlantic/Madeira:00:00	Atlantic/Stanley:-:04:00
Atlantic/Cape_Verde:-:01:00	Atlantic/Reykjavik:00:00	—

Table B-12 Australia GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Australia/ACT:+:10:00	Australia/LHI:+:10:30	Australia/Queensland:+:10:00
Australia/Adelaide:+:09:30	Australia/Lindeman:+:10:00	Australia/South:+:09:30
Australia/Brisbane:+:10:00	Australia/Lord_Howe:+:10:30	Australia/Sydney:+:10:00
Australia/Broken_Hill:+:09:30	Australia/Melbourne:+:10:00	Australia/Tasmania:+:10:00
Australia/Canberra:+:10:00	Australia/NSW:+:10:00	Australia/Victoria:+:10:00
Australia/Darwin:+:09:30	Australia/North:+:09:30	Australia/West:+:08:00
Australia/Hobart:+:10:00	Australia/Perth:+:08:00	Australia/Yancowinna:+:09:30

Table B-13 Brazil GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Brazil/Acre:-:05:00	Brazil/East:-:03:00	Brazil/West:-:04:00
Brazil/DeNoronha:-:02:00	—	—

Table B-14 Canada/Chile/Cuba GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Canada/Atlantic:-:04:00	Canada/Mountain:-:07:00	Canada/Yukon:-:08:00
Canada/Central:-:06:00	Canada/Newfoundland:-:03:30	Chile/Continental:-:04:00
Canada/East-Ssktchwan:-:06:00	Canada/Pacific:-:08:00	Chile/EasterIsland:-:06:00
Canada/Eastern:-:05:00	Canada/Saskatchewan:-:06:00	Cuba:-:05:00

Table B-15 Egypt/Eire/Europe GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Egypt:+:02:00	Europe/Kiev:+:02:00	Europe/Simferopol:+:02:00
Eire:00:00	Europe/Lisbon:00:00	Europe/Skopje:+:01:00
Europe/Amsterdam:+:01:00	Europe/Ljubljana:+:01:00	Europe/Sofia:+:02:00
Europe/Andorra:+:01:00	Europe/London:00:00	Europe/Stockholm:+:01:00

Table B-15 Egypt/Eire/Europe GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Europe/Athens:+:02:00	Europe/Luxembourg:+:01:00	Europe/Tallinn:+:02:00
Europe/Belfast:00:00	Europe/Madrid:+:01:00	Europe/Tirane:+:01:00
Europe/Belgrade:+:01:00	Europe/Malta:+:01:00	Europe/Tiraspol:+:02:00
Europe/Berlin:+:01:00	Europe/Minsk:+:02:00	Europe/Uzhgorod:+:02:00
Europe/Bratislava:+:01:00	Europe/Monaco:+:01:00	Europe/Vaduz:+:01:00
Europe/Brussels:+:01:00	Europe/Moscow:+:03:00	Europe/Vatican:+:01:00
Europe/Bucharest:+:02:00	Europe/Nicosia:+:02:00	Europe/Vienna:+:01:00
Europe/Budapest:+:01:00	Europe/Oslo:+:01:00	Europe/Vilnius:+:02:00
Europe/Chisinau:+:02:00	Europe/Paris:+:01:00	Europe/Warsaw:+:01:00
Europe/Copenhagen:+:01:00	Europe/Prague:+:01:00	Europe/Zagreb:+:01:00
Europe/Dublin:00:00	Europe/Riga:+:02:00	Europe/Zaporozhye:+:02:00
Europe/Gibraltar:+:01:00	Europe/Rome:+:01:00	Europe/Zurich:+:01:00
Europe/Helsinki:+:02:00	Europe/Samara:+:04:00	Europe/Simferopol:+:02:00
Europe/Istanbul:+:02:00	Europe/San_Marino:+:01:00	Europe/Skopje:+:01:00
Europe/Kaliningrad:+:02:00	Europe/Sarajevo:+:01:00	Europe/Sofia:+:02:00

Table B-16 Hong Kong/Iceland/India/Iran/Israel GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Hongkong:+:08:00	Indian/Cocos:+:06:30	Indian/Mauritius:+:04:00
Iceland:00:00	Indian/Comoro:+:03:00	Indian/Mayotte:+:03:00
Indian/Antananarivo:+:03:00	Indian/Kerguelen:+:05:00	Indian/Reunion:+:04:00
Indian/Chagos:+:06:00	Indian/Mahe:+:04:00	Iran:+:03:30
Indian/Christmas:+:07:00	Indian/Maldives:+:05:00	Israel:+:02:00

Table B-17 Jamaica/Japan/Kwajalein/Libya GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Jamaica:-:05:00	Kwajalein:+:12:00	Libya:+:02:00
Japan:+:09:00	—	—

Table B-18 Mexico/Mideast GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Mexico/BajaNorte:-:08:00	Mexico/General:-:06:00	Mideast/Riyadh88:+:03:07
Mexico/BajaSur:-:07:00	Mideast/Riyadh87:+:03:07	Mideast/Riyadh89:+:03:07

Table B-19 Pacific/Poland/Portugal GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Pacific/Apia:-:11:00	Pacific/Johnston:-:10:00	Pacific/Ponape:+:11:00
Pacific/Auckland:+:12:00	Pacific/Kiritimati:+:14:00	Pacific/Port_Moresby:+:10:00
Pacific/Chatham:+:12:45	Pacific/Kosrae:+:11:00	Pacific/Rarotonga:-:10:00
Pacific/Easter:-:06:00	Pacific/Kwajalein:+:12:00	Pacific/Saipan:+:10:00
Pacific/Efate:+:11:00	Pacific/Majuro:+:12:00	Pacific/Samoa:-:11:00
Pacific/Enderbury:+:13:00	Pacific/Marquesas:-:09:30	Pacific/Tahiti:-:10:00
Pacific/Fakaofo:-:10:00	Pacific/Midway:-:11:00	Pacific/Tarawa:+:12:00
Pacific/Fiji:+:12:00	Pacific/Nauru:+:12:00	Pacific/Tongatapu:+:13:00
Pacific/Funafuti:+:12:00	Pacific/Niue:-:11:00	Pacific/Truk:+:10:00
Pacific/Galapagos:-:06:00	Pacific/Norfolk:+:11:30	Pacific/Wake:+:12:00
Pacific/Gambier:-:09:00	Pacific/Noumea:+:11:00	Pacific/Wallis:+:12:00
Pacific/Guadalcanal:+:11:00	Pacific/Pago_Pago:-:11:00	Pacific/Yap:+:10:00
Pacific/Guam:+:10:00	Pacific/Palau:+:09:00	Poland:+:01:00
Pacific/Honolulu:-:10:00	Pacific/Pitcairn:-:08:00	Portugal:00:00

Table B-20 Singapore/System V/Turkey GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Singapore:+:08:00	SystemV/EST5:-:05:00	SystemV/PST8PDT:-:08:00
SystemV/AST4:-:04:00	SystemV/EST5EDT:-:05:00	SystemV/YST9:-:09:00
SystemV/AST4ADT:-:04:00	SystemV/MST7:-:07:00	SystemV/YST9YDT:-:09:00
SystemV/CST6:-:06:00	SystemV/MST7MDT:-:07:00	Turkey:+:02:00
SystemV/CST6CDT:-:06:00	SystemV/PST8:-:08:00	—

Table B-21 U.S. GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
US/Alaska:-:09:00	US/Eastern:-:05:00	US/Pacific-New:-:08:00
US/Aleutian:-:10:00	US/Hawaii:-:10:00	US/Pacific:-:08:00
US/Arizona:-:07:00	US/Indiana-Starke:-:05:00	US/Samoa:-:11:00
US/Central:-:06:00	US/Michigan:-:05:00	—
US/East-Indiana:-:05:00	US/Mountain:-:07:00	—

APPENDIX C

Creating Coverage Zone Files

This appendix describes the Coverage Zone file and provides several Coverage Zone file examples.

Introduction

A *Coverage Zone file* is an XML file used to specify a user-defined coverage zone. The Coverage Zone file supports different tags to support different types of proximity configurations.

- Network and subnet—Specify the IP address range
- Geographical location—Specify the longitude and latitude of the data center

In addition to the coverage zone information, two optional elements are created for documentation purposes: a revision value to specify the version of the Coverage Zone file and a customer name.

For information about importing or uploading a Coverage Zone file, see the "Coverage Zone File Registration" section on page 6-9.

For more information about Coverage Zone files, see the "Coverage Zone File" section on page 1-21.

Coverage Zone files can be created using any ASCII text-editing tool. You can use a single coverage zone text-format file to define all the coverage zones for your CDS network.

Table C-1 defines the Coverage Zone file elements.

Table C-1 Coverage Zone File Elements

Tag	Element	Value	Description
location	latitude	float	Value indicating the geographical coordinate (latitude) of the data center.
	longitude	float	Value indicating the geographical coordinate (longitude) of the data center.
coverageZone	network	IP address	Coverage zone IP address range.
	SE	Service Engine name (string)	Specifies the Service Engines serving the coverage zone specified in the network element. This can have one or more elements.
	metric	integer	Value indicating the proximity of the Service Engine to the end user. The lower the value, the closer the Service Engine is to the end user.
	location		Value indicating the geographical coordinates (latitude and longitude) of the data center.

Table C-1 Coverage Zone File Elements (continued)

Tag	Element	Value	Description
CDNNetwork	revision	1.0	Not used in this CDS release.
	customerName	customer name	Not used in this CDS release.
	coverageZone		This can have one or more coverage zones.

Note The metric value of a default coverage zone is set to 20. If a particular SE is preferred for a user-defined coverage zone, the metric value in the Coverage Zone file should be set to a value less than 20. If a default coverage zone is preferred, then the metric value in the Coverage Zone file should be set to a value greater than 20.

Coverage Zone XML Schema

The XML Schema file describes and dictates the content of the XML file. The coveragezone.xsd file contains the XML schema.

The following code is the Coverage Zone XML schema:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema">
    <xsd:annotation>
        <xsd:documentation>Revision: 1.0</xsd:documentation>
    </xsd:annotation>
    <xsd:annotation>
        <xsd:documentation> Schema used to validate Cisco CoverageZone file </xsd:documentation>
    </xsd:annotation>
    <xsd:simpleType name="ipAddressType">
        <!-- IPv4 Address -->
        <!-- Snum = 1*3Digit ; representing a decimal integer -->
        <!--                 ; value in the range 0 through 255 -->
        <!-- IPv4-address-literal - Snum 3("." Snum) -->
        <xsd:restriction base="xsd:token">
            <xsd:pattern value="((1?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])\.){3}(1?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])/[0-9]+"/>
        </xsd:restriction>
    </xsd:simpleType>
    <xsd:simpleType name="SENameType">
        <xsd:restriction base="xsd:string">
            <xsd:maxLength value="80"/>
        </xsd:restriction>
    </xsd:simpleType>
    <xsd:simpleType name="LatitudeType">
        <xsd:restriction base="xsd:float">
            <xsd:minInclusive value="90" />
            <xsd:maxInclusive value="90" />
        </xsd:restriction>
    </xsd:simpleType>
    <xsd:simpleType name="LongitudeType">
        <xsd:restriction base="xsd:float">
            <xsd:minInclusive value="180" />
            <xsd:maxInclusive value="180" />
        </xsd:restriction>
    </xsd:simpleType>
    <xsd:complexType name="locationType">
        <xsd:sequence>
```

```
            <xsd:element name="latitude" type="LatitudeType">
            <xsd:element name="longitude" type="LongitudeType">
        </xsd:sequence>
    </xsd:complexType>
    <xsd:complexType name="coverageZoneType">
        <xsd:sequence>
            <xsd:choice maxOccurs="1">
                <xsd:sequence maxOccurs="2">
                    <xsd:element name="network" type="ipAddressType"/>
                    <xsd:element name="location" type="locationType"/>
                </xsd:sequence>
            <xsd:choice maxOccurs="1">
                    <xsd:element name="network" type="ipAddressType"/>
                    <xsd:element name="location" type="locationType"/>
            </xsd:choice>
            </xsd:choice>
            <xsd:element name="SE" type="SENameType" maxOccurs="unbounded" />
            <xsd:element name="metric" type="xsd:positiveInteger"maxOccurs="1" />
        </xsd:sequence>
    </xsd:complexType>
    <xsd:element name="CDNNetwork">
        <xsd:complexType>
            <xsd:sequence>
                <xsd:element name="revision" type="xsd:string" minOccurs="0" maxOccurs="1" />
                <xsd:element name="customerName" type="xsd:string" minOccurs="0" maxOccurs="1" />
                <xsd:element name="coverageZone" type="coverageZoneType" minOccurs="0" maxOccurs="unbounded"/>
            </xsd:sequence>
        </xsd:complexType>
    </xsd:element>
</xsd:schema>
```

Coverage Zone File Examples

The following sections show different Coverage Zone file examples in three scenarios.

Scenario 1: Coverage Zone with Client Network Only

```
<?xml version=”1.0” ?>
<!-- Coverage Zone data in XML -->
<CDNNetwork>
    <revision>1.0</revision>
    <customerName>Cisco Systems</customerName>
 <!-- San Jose Datacenter -->
    <coverageZone>
        <network> 192.1.2.0/16 </network>
        <SE> CDE-200-SE1 </SE>
        <SE> CDE-200-SE2 </SE>
        <metric> 10 </metric>
    </coverageZone>
<!-- Chicago Datacenter -->
    <coverageZone>
        <network> 192.1.3.0/24 </network>
        <SE> CDE-200-SE3 </SE>
        <SE> CDE-200-SE4 </SE>
        <metric> 10 </metric>
    </coverageZone>
<!-- New York Datacenter -->
    <coverageZone>
        <network> 192.1.4.0/24 </network>
```

```
        <SE> CDE-200-SE5 </SE>
        <SE> CDE-200-SE6 </SE>
        <metric> 10 </metric>
    </coverageZone>
</CDNNetwork>
```

Scenario 2: Coverage Zone with Geographical Location of the Datacenter Only

```
<?xml version="1.0" ?>
<!-- Coverage Zone data in XML -->
<CDNNetwork>
    <revision>1.0</revision>
    <customerName>Cisco Systems</customerName>
<!-- San Jose Datacenter -->
    <coverageZone>
        <location>
            <latitude> 20 </latitude>
            <longitude> -120 </longitude>
        </location>
        <SE> CDE-200-SE1 </SE>
        <SE> CDE-200-SE2 </SE>
        <metric> 10 </metric>
    </coverageZone>
<!-- Chicago Datacenter -->
    <coverageZone>
        <location>
            <latitude> 20 </latitude>
            <longitude> -120 </longitude>
        </location>
        <SE> CDE-200-SE3 </SE>
        <SE> CDE-200-SE4 </SE>
        <metric> 10 </metric>
    </coverageZone>
<!-- New York Datacenter -->
    <coverageZone>
        <location>
        <latitude> 20 </latitude>
            <longitude> -120 </longitude>
            </location>
        <SE> CDE-200-SE5 </SE>
        <SE> CDE-200-SE6 </SE>
        <metric> 10 </metric>
    </coverageZone>
</CDNNetwork>
```

Scenario 3: Coverage Zone with Client Network and Geographical Location of the Datacenter

```
<?xml version="1.0" ?>
<!-- Coverage Zone data in XML -->
<CDNNetwork>
    <revision>1.0</revision>
    <customerName>Cisco</customerName>
<!-- San Jose Datacenter -->
    <coverageZone>
        <network> 192.1.2.0/16 </network>
        <SE> CDE-200-SE1 </SE>
        <SE> CDE-200-SE2 </SE>
        <metric> 10 </metric>
```

```
    </coverageZone>
 <!-- Chicago Datacenter -->
    <coverageZone>
        <location>
            <latitude> 20 </latitude>
            <longitude> -120 </longitude>
        </location>
        <SE> CDE-200-SE3 </SE>
        <SE> CDE-200-SE4 </SE>
        <metric> 10 </metric>
    </coverageZone>
<!-- New York Datacenter -->
    <coverageZone>
        <network> 192.1.4.0/24 </network>
        <SE> CDE-200-SE5 </SE>
        <SE> CDE-200-SE6 </SE>
        <metric> 10 </metric>
    </coverageZone>
</CDNNetwork>
```

APPENDIX D

Creating and Manipulating Session Shifting Files

This appendix describes the process of creating Session Shifting files, and includes the following topics:

- Introduction, page D-1
- Content Manager, page D-2
- Subscriber Manager, page D-4
- Profile Manager, page D-5
- Session Manager, page D-7

Note 3-Screen Session Shifting is a Release 2.2 feature and supports RTSP streaming for Windows Media Streaming and the Movie Streamer.

Introduction

Note Release 2.2 supports web-based programming APIs. These APIs are based on the REpresentational State Transfer (REST) architecture. In Release 2.2, 3-Screen Session Shifting and Capability Exchange are the only applications that are supported through the Web Services infrastructure.

The Internet Streamer CDS incorporates a Web Services infrastructure in order to support the 3-Screen Session Shifting feature. The Web Services interface provides a way for service providers to communicate information about content and subscribers to the CDS for session shifting.

The Web Services infrastructure features that are supported in this release are the following:

- Content Manager
- Subscriber Manager
- Profile Manager
- Stream Session Event Manager

The Service Router, designated as the *centralized session manager,* manages each Web Services manager. Each Web Services manager has an XML file and associated HTTP GET message formats that use the REST architecture. The XML files for the Content Manager, Subscriber Manager, and Profile Manager are imported through the delivery service Session Shifting page. The Session Shifting page offers the ability to add, append, and delete these files. See the "Session Shifting" section on page 5-10 for more information.

Session Shifting and Delivery Services

The CDS uses the delivery services to implement most of its functions. Each delivery service is associated with an origin server. A Content Origin in the CDS defines a Service Routing Domain Name, which is used by the clients to receive the content delivery from the CDS platform. The Service Routing Domain Name is only used for serving content. The Service Routing Domain Name configured for the Content Origin should also be configured in the DNS servers, so client requests can get redirected to a Service Router for request mediation and redirection.

The Web Services interface for session shifting is implemented on a per-delivery service basis. A Service Router, identified by the Service Routing Domain Name, is designated as the *centralized session manager* for each delivery service participating in session shifting. Only one Service Router is used for all session shifting delivery services. A delivery service participating in session shifting can also be used to deliver content.

Note The schemas used in the CDS for the Web Services interface follow the SCTE-130 standard and are available upon request.

Content Manager

The Content Manager facilitates the management of content information for 3-Screen Session Shifting. Content Lists can be posted to a Web Services interface in order to add or delete content. The provisioned content is visible at a simple URL.

For example, for the delivery service with a Service Routing Domain Name of sr.cds.com, the following link can be used to address the Content Manager to provision content:

```
http://sr.cds.com/api/services/content_mgr
```

Content Lists

Content Lists in XML format are used to add, append, and delete content entries through the delivery service Session Shifting page. The following is an example of a Content List for a delivery service with a Service Routing Domain Name of cds.com:

```
<?xml version="1.0"?>
<ws:ContentList

    xmlns="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:ws="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:core="http://www.scte.org/schemas/130-2/2007/core"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation=
    "http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc ciscowebsvc.xsd''>

<core:Content>
<core:URI>rtsp://cds.com/pirates.mov</core:URI>
    <core:Ext>
        <core:Content>
            <core:URI>rtsp://cds.com/pirates.pc.mov</core:URI>
        </core:Content>
        <core:Content>
```

```
            <core:URI>rtsp://cds.com/pirates.tv.mov</core:URI>
        </core:Content>
        <core:Content>
            <core:URI>rtsp://cds.com/pirates.pda.wmv</core:URI>
        </core:Content>
        <core:Content>
            <core:AssetRef providerID="disney.com"
                           assetID="aaaa0000000000000001"/>
        </core:Content>
    </core:Ext>
</core:Content>

<core:Content>
<core:URI>rtsp://cds.com/lost.mov</core:URI>
    <core:Ext>
        <core:Content>
            <core:URI>rtsp://cds.com/lost.pc.mov</core:URI>
        <core:Content>
            <core:URI>rtsp://cds.com/lost.tv.mov</core:URI>
        </core:Content>
        <core:Content>
            <core:URI>rtsp://cds.com/lost.pda.wmv</core:URI>
        <core:Content>
            <core:AssetRef providerID="abc.com"
                           assetID="aaaa0000000000000002"/>
        </core:Content>
    </core:Ext>
</core:Content>

</ws:ContentList>
```

This Content List has two sets of content objects, one for “Pirates” and one for “Lost.” Their grouping indicates that all the listed URIs and AssetRefs within each set has content that is visually similar, while each might be formatted for a specific client device. For example, the movie called “Pirates,” has separate URIs for each of the following client devices, PC, TV, and PDA. The filename and filename termination syntax used in this example are for illustrative purposes; you can use any name that is appropriate.

Supported Content Management Operations

The information content sets of a Content List can be added, appended, deleted, or retrieved to and from the centralized session manager through HTTP messages. The creation of the HTTP message itself depends on the client.

For illustrative purposes, the following examples use the **wget** command, using the file content_list.xml.

To add or append a Content List, use the following HTTP message:

```
wget --post-file=content_list.xml http://cds.com/api/services/content_mgr/
```

To delete a Content List, use the following HTTP message:

```
wget --post-file=content_list.xml
http://cds.com/api/services/content_mgr/\(_method=DELETE\)
```

Note The backslashes preceding the opening and closing parentheses are escape characters for **wget**. If a different client is used, the characters required may be different.

To retrieve a Content List with the contents of the database, use the following **wget** command:

```
wget http://cds.com/api/services/content_mgr/
```

Note This URL can also be called from a web browser such as Internet Explorer.

Subscriber Manager

The Subscriber Manager facilitates the management of 3-Screen Session Shifting users. Subscriber Lists can be posted to a Web Services interface in order to add or delete subscribers. The provisioned subscribers are visible at a simple URL.

For example, for the delivery service with a Service Routing Domain Name of sr.cds.com, the following link can be used to address the Subscriber Manager to provision subscribers:

```
http://sr.cds.com/api/services/subscriber_mgr
```

Subscriber Lists

Subscriber Lists in XML format are used to add, append, and delete content entries through the Subscriber Manager. The following is an example of a Subscriber List for a delivery service:

```
<?xml version="1.0"?>
<ws:SubscriberList

    xmlns="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:ws ="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:sisdb="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscoTadmSIS"
    xmlns:adm="http://www.scte.org/schemas/130-3/2007/adm"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation=
    "http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc ciscowebsvc.xsd''>

<sisdb:BillableEntity id="SmithFamily">
    <adm:TerminalAddress type="STB-MAC">02:05:05:74:AA:EE</adm:TerminalAddress>
    <sisdb:User id="BobSmith">
        <adm:TerminalAddress type="LOGIN">bob</adm:TerminalAddress>
    </sisdb:User>
    <sisdb:User id="BeckySmith">
        <adm:TerminalAddress type="LOGIN">becky</adm:TerminalAddress>
    </sisdb:User>
    <sisdb:User id="JohnSmith">
        <adm:TerminalAddress type="LOGIN">john</adm:TerminalAddress>
    </sisdb:User>
</sisdb:BillableEntity>

<sisdb:BillableEntity id="JonesFamily">
    <adm:TerminalAddress type="STB-MAC">02:05:05:74:AA:EE</adm:TerminalAddress>
    <sisdb:User id="TimJones">
        <adm:TerminalAddress type="LOGIN">tim</adm:TerminalAddress>
    </sisdb:User>
    <sisdb:User id="AliceJones">
        <adm:TerminalAddress type="LOGIN">alice</adm:TerminalAddress>
    </sisdb:User>
    <sisdb:User id="KateJones">
```

```
        <adm:TerminalAddress type="LOGIN">kate</adm:TerminalAddress>
    </sisdb:User>
</sisdb:BillableEntity>

</ws:SubscriberList>
```

This Subscriber List has two sets of subscribers, one for the Smith Family and one for the Jones Family. Their grouping indicates that all the listed terminal addresses may correspond to the same viewers group. For example, Bob Smith might be watching certain content from Terminal Address STB-MAC = 02:05:05:74:AA:EE, pause, and afterwards want to continue watching the same content from his PC using a LOGIN = bob. The grouping into a single BillableEntity makes it possible for CDS to maintain the session for Bob Smith regardless of which of the two clients he is using (TV with a particular STB-MAC or a PC with a particular LOGIN).

Note The terminal address type and value can be any strings that are appropriate.

Supported Subscriber Management Operations

The information subscriber sets of a Subscriber List can be added, appended, deleted, or retrieved to and from the centralized session manager through HTTP messages. The creation of the HTTP message itself depends on the client.

For illustrative purposes, the following examples use the **wget** command, using the file subscriber_list.xml.

To add or append a Subscriber List, use the following HTTP message:

```
wget --post-file=subscriber_list.xml http://cds.com/api/services/subscriber_mgr/
```

To delete a Subscriber List, use the following HTTP message:

```
wget --post-file=subscriber_list.xml
http://cds.com/api/services/subscriber_mgr/\(_method=DELETE\)
```

Note The backslashes preceding the opening and closing parentheses are escape characters for **wget**. If a different client is used, the characters required may be different.

To retrieve a Subscriber List with the contents of the database, use the following **wget** command:

```
wget http://cds.com/api/services/subscriber_mgr/
```

Note This URL can also be called from a web browser such as Internet Explorer.

Profile Manager

The Profile Manager facilitates the management of 3-Screen Session Shifting profiles for Capability Exchange. Capability Exchange refers to the process of mapping the logical URL to a specific content format inside the group. One profile groups together multiple clients that have similar characteristics, such as supported file format, desired resolution, or bandwidth. Profile Lists can be posted to a Web Services interface in order to add or delete profiles. The provisioned profile is visible at a simple URL.

For example, for the delivery service with a Service Routing Domain Name of sr.cds.com, the following link can be used to address the Profile Manager to provision profiles:

```
http://sr.cds.com/api/services/profile_mgr
```

Profile Lists

Profile Lists in XML format are used to add, append, and delete profile entries through the Profile Manager. The following is an example of a Profile List for a delivery service:

```
<?xml version="1.0"?>
<ws:ProfileList

    xmlns="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:ws="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation=
    "http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc ciscowebsvc.xsd''>

<Profile id="1">
    <UserAgent>QuickTime</UserAgent>
    <CapabilitiesUrl>http://apple.com/capabilities/quicktime</CapabilitiesUrl>
</Profile>

<Profile id="2">
    <UserAgent>WMPlayer</UserAgent>
    <CapabilitiesUrl>http://microsoft.com/capabilities/wmplayer</CapabilitiesUrl>
</Profile>

</ws:ProfileList
```

This Profile List has two profiles defined. The first one will be used for QuickTime clients that identify themselves as a QuickTime user agent or provide the capability URL shown. The second one will be used for Windows Media Player clients that identify themselves as a Windows Media Player user agent or provide the capabilities URL shown.

Note The user agent and capabilities URL depend on the URL provided by the client in the RTSP messages. Also note that the example above creates a single profile for all QuickTime clients, regardless of their supported bandwidth or other specifications. If it is required to have profiles defined with finer granularities, this can be done as long as the profiles have different user agents and capabilities URLs.

Profiles in Content Lists

When provisioning content entries using Content Lists, a profile ID can be specified for each content entry within a content set. This is done by including a Profile tag in the Ext element of the particular content entry. If no profile ID is specified for a particular entry, a default of Profile id= “0” is given. The following content set has been modified to include profile IDs:

```
<core:Content>
<core:URI>rtsp://cds.com/pirates.mov</core:URI>
    <core:Ext>
        <Profile id=”1”>
        <core:Content>
            <core:URI>rtsp://cds.com/pirates.pc.mov</core:URI>
```

```
        </core:Content>
        <core:Content>
            <core:URI>rtsp://cds.com/pirates.tv.mov</core:URI>
        </core:Content>
        <core:Content>
            <core:URI>rtsp://cds.com/pirates.pda.wmv</core:URI>
            <core:Ext>
                <Profile id=”2”>
            <core:Ext>
        </core:Content>
        <core:Content>
            <core:AssetRef providerID="disney.com"
                           assetID="aaaa0000000000000001"/>
        </core:Content>
    </core:Ext>
</core:Content>
```

In this Content List, rtsp://cds.com/pirates.mov will be entered with Profile id=“1,” and rtsp://cds.com/pirates.pda.wmv will be entered with Profile id=“2.” The other entries will have a default Profile id=“0.”

Supported Profile Management Operations

The profile information of a Profile List can be added, appended, deleted, or retrieved to and from the centralized session manager through HTTP messages. The creation of the HTTP message itself depends on the client.

For illustrative purposes, the following examples use the **wget** command, using the file profile_list.xml.

To add or append a Profile List, use the following HTTP message:

```
wget --post-file=profile_list.xml http://cds.com/api/services/profile_mgr/
```

To delete a Profile List, use the following HTTP message:

```
wget --post-file=profile_list.xml
http://cds.com/api/services/profile_mgr/\(_method=DELETE\)
```

Note The backslashes preceding the opening and closing parentheses are escape characters for **wget**. If a different client is used, the characters required may be different.

To retrieve a Profile List with the profiles of the database, use the following **wget** command:

```
wget http://cds.com/api/services/profile_mgr/
```

Note This URL can also be called from a web browser such as Internet Explorer.

Session Manager

The Stream Session Event Manager allows creating and retrieving “last play times” for pairs of content and subscribers. The session information is stored by content set and subscriber set. Thus, in the examples illustrated in the “Content Lists” section on page D-2, and the “Subscriber Lists” section on page D-4, Bob Smith might have started watching rtsp://cds.com/pirates.mov on his set-top box with

STB-MAC = 02:05:05:74:AA:EE, paused, and then continued watching from his PC with LOGIN=bob from the URL rtsp://cds.com/pirates.pc.mov. The movie restarts from where he last paused it, even though the URL is different (though it is in the same content set) and the terminal address is different (though it is in the same BillableEntity).

Stream Session Event

A Stream Session Event specifies a pause or stop time for a particular content and a particular terminal address. The following is an example of a Stream Session Event:

```
<?xml version="1.0"?>
<ws:StreamSessionEvent

    xmlns="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:ws ="http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc"
    xmlns:core="http://www.scte.org/schemas/130-2/2007/core"
    xmlns:adm="http://www.scte.org/schemas/130-3/2007/adm"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation=
    "http://www.cisco.com/schemas/VCNBU/SCTE130Impl/2007/R1/ciscowebsvc ciscowebsvc.xsd''>

<core:Content>
    <core:URI>rtsp://cds.com/pirates.mov</core:URI>
</core:Content>

<adm:NPT scale="1.0">600</adm:NPT>

<adm:Client>
    <adm:TerminalAddress type="LOGIN">bob</adm:TerminalAddress>
</adm:Client>

</ws:StreamSessionEvent>
```

This particular Stream Session Event specifies that the content in URL rtsp://cds.com/pirates.mov was paused 600 seconds after the start time (10 minutes into the video) by the user with LOGIN=bob.

Session Update

The session last play time information can be updated or looked up from the centralized session manager through HTTP messages. The creation of the HTTP message itself depends on the client.

For illustrative purposes, the following examples use the **wget** command, using the file session_event.xml.

To update a session, use the following **wget** command:

```
wget --post-file=session_event.xml http://cds.com/api/services/session_event_mgr/
```

Note The session_event.xml file has a Stream Session Event as described in the "Stream Session Event" section.

Session Lookup

To look up the last play time, you need to provide the content and terminal address for which you are querying the centralized session manager. Because the GET method does not support providing information to the centralized session manager, you must use an overloaded POST message with the GET method as shown in the following example:

```
wget --post-file=session_lookup.xml
http://cds.com/api/services/session_event_mgr/\(_method=GET\)
```

The file displayed with this HTTP message has the last play time in the normal play time (NPT) value.

Note The backslashes preceding the opening and closing parentheses are escape characters for **wget**. If a different client is used, the characters required may be different.

APPENDIX E

CLI Commands

This appendix covers the following topics:

- Configuring Port Channel, page E-1
- Configuring Last-Resort Routing, page E-6
- Configuring Standby Interfaces, page E-6
- Other CLI Commands, page E-10

Configuring Port Channel

To configure an EtherChannel, you use the **PortChannel** interface configuration command. Port Channel, also known as EtherChannel, supports the grouping of up to four same-speed network interfaces into one virtual interface. This grouping capability allows the setting or removing of a virtual interface that consists of two, three, or four Gigabit Ethernet interfaces. EtherChannel also provides interoperability with Cisco routers, switches, and other networking devices or hosts supporting EtherChannel; load balancing; and automatic failure detection and recovery based on each interface's current link status.

Note In order to achieve the best throughput, we recommend you configure a port channel for the four Gigabit Ethernet ports on the line card.

Redundant Dedicated Management Ports

On a CDE200 configured as an SE or SR there are six Gigabit Ethernet ports. All of the ports can be used for delivery traffic such as RTSP, as well as system management traffic to communicate with other CDS devices such as the CDSM. In order to prevent all the bandwidth being used by delivery traffic, a dedicated management port setup is often recommended.

In case of physical failure on a single port, channel bonding configuration of multiple Gigabit Ethernet ports is also recommended for both delivery traffic and management traffic.

Note A port channel configured with a default gateway is only for delivery traffic. Delivery traffic places highest bandwidth demand on the CDS network. A port channel configured as the primary interface carries delivery traffic.

Note If an EtherChannel (also known as port channel) is used between the upstream router or switch and the SE for streaming real-time data, the EtherChannel load balance algorithms on the upstream switch or router and the SE should be configured as "src-ip" and "dst-ip" respectively. Using this configuration ensures session stickiness and general balanced load distribution based on clients' IP addresses. Also, distribute your client IP address space across multiple subnets so that the load balancing algorithm is effective in spreading the traffic among multiple ports.

Typically, two Gigabit Ethernet ports on the motherboard (GigabitEthernet 1/0 and 2/0) can be bundled for the management port channel, and four Gigabit Ethernet ports (GigabitEthernet 3/0, 4/0, 5/0, and 6/0) on the network interface card (NIC) can be bundled for the traffic port channel.

To configure redundant dedicated management ports using the CLI, do the following:

Step 1 Configure two port channels., one with four Gigabit Ethernet ports (portchannel 1) and one with two Gigabit Ethernet ports (portchannel 2) for application traffic and management traffic respectively.

```
SE(config)# interface GigabitEthernet 1/0
SE(config-if)# channel-group 2
SE(config-if)# exit
SE(config)# interface GigabitEthernet 2/0
SE(config-if)# channel-group 2
SE(config-if)# exit
SE(config)# interface GigabitEthernet 3/0
SE(config-if)# channel-group 1
SE(config-if)# exit
SE(config)# interface GigabitEthernet 4/0
SE(config-if)# channel-group 1
SE(config-if)# exit
SE(config)# interface GigabitEthernet 5/0
SE(config-if)# channel-group 1
SE(config-if)# exit
SE(config)# interface GigabitEthernet 6/0
SE(config-if)# channel-group 1
SE(config-if)# exit
```

Note The port channel carrying delivery traffic should always be configured as channel-group 1 and set as the primary interface.

Step 2 Configure two different subnets for both port channels.

```
SE(config)# interface PortChannel 1
SE(config-if)# ip address 3.1.7.73 255.255.255.0
SE(config-if)# exit
SE(config)# interface PortChannel 2
SE(config-if)# ip address 3.1.8.200 255.255.255.0
SE(config-if)# exit
```

Step 3 Configure the delivery port channel as the primary interface.

```
SE(config)# primary-interface PortChannel 1
```

Step 4 Configure a default gateway for the delivery traffic.

```
SE(config)# ip default-gateway 3.1.7.1
```

Step 5 Set the load balancing algorithm to the destination IP address.

```
SE(config)# port-channel load-balance dst-ip
```

Step 6 Set the client and server TCP compliance to the RFC 1323 standard.

```
SE(config)# tcp client-satellite
SE(config)# tcp server-satellite
```

Step 7 Configure a static route to the CDSM (4.0.5.5) to specify that all management traffic will go through this interface.

```
SE(config)# ip route 4.0.5.5 255.255.255.255 3.1.8.1
```

Step 8 Configure the port channel and VLANs on the switch that the SE is directly connected to.

```
SW3750(config)# interface Port-channel1
SW3750(config-if)# switchport access vlan 201
SW3750(config-if)# exit
SW3750(config)# interface Port-channel2
SW3750(config-if)# switchport access vlan 202
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/1
SW3750(config-if)# description Connected to portchannel2
SW3750(config-if)# switchport access vlan 202
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/2
SW3750(config-if)# description Connected to portchannel2
SW3750(config-if)# switchport access vlan 202
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/3
SW3750(config-if)# description connected to portchannel1
SW3750(config-if)# switchport access vlan 201
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/4
SW3750(config)-if# description connected to portchannel1
SW3750(config-if)# switchport access vlan 201
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/5
SW3750(config-if)# description connected to portchannel1
SW3750(config-if)# switchport access vlan 201
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface GigabitEthernet1/0/6
SW3750(config-if)# description connected to portchannel1
SW3750(config-if)# switchport access vlan 201
SW3750(config-if)# switchport mode access
SW3750(config-if)# exit
SW3750(config)# interface Vlan201
SW3750(config-if)# ip address 3.1.7.1 255.255.255.0
SW3750(config-if)# exit
SW3750(config)# interface Vlan202
SW3750(config-if)# ip address 3.1.8.1 255.255.255.0
SW3750(config-if)# exit
```

Step 9 Set the load balancing algorithm to the source IP address.

```
SW3750(config)# port-channel load-balance src-ip
```

To verify the setup before application traffic is sent, use the following:

```
SE# clear statistics all
SE# show interface portChannel 1
Interface PortChannel 1 (2 physical interface(s)):
GigabitEthernet 3/0 (active)
GigabitEthernet 4/0 (active)
GigabitEthernet 5/0 (active)
GigabitEthernet 6/0 (active)
---------------------
Type:Ethernet
Ethernet address:00:04:23:D8:86:02
Internet address:3.1.7.73
Broadcast address:3.1.7.255
Netmask:255.255.255.0
Maximum Transfer Unit Size:1500
Metric:1
Packets Received: 28
Input Errors: 0
Input Packets Dropped: 0
Input Packets Overruns: 0
Input Packets Frames: 0
Packet Sent: 40
Output Errors: 0
Output Packets Dropped: 0
Output Packets Overruns: 0
Output Packets Carrier: 0
Output Queue Length:0
Collisions: 0
Flags:UP BROADCAST RUNNING MASTER MULTICAST

SE# show interface portChannel 2
Interface PortChannel 2 (4 physical interface(s)):
GigabitEthernet 1/0 (active)
GigabitEthernet 2/0 (active)
---------------------
Type:Ethernet
Ethernet address:00:30:48:33:01:26
Internet address:3.1.8.200
Broadcast address:3.1.8.255
Netmask:255.255.255.0
Maximum Transfer Unit Size:1500
Metric:1
Packets Received: 6
Input Errors: 0
Input Packets Dropped: 0
Input Packets Overruns: 0
Input Packets Frames: 0
Packet Sent: 0
Output Errors: 0
Output Packets Dropped: 0
Output Packets Overruns: 0
Output Packets Carrier: 0
Output Queue Length:0
Collisions: 0
Flags:UP BROADCAST RUNNING MASTER MULTICAST
```

To verify the setup after application traffic is sent, use the following:

```
SE# show interface portChannel 1
Interface PortChannel 1 (4 physical interface(s)):
GigabitEthernet 3/0 (active)
GigabitEthernet 4/0 (active)
GigabitEthernet 5/0 (active)
```

```
GigabitEthernet 6/0 (active)
--------------------
Type:Ethernet
Ethernet address:00:04:23:D8:86:02
Internet address:3.1.7.73
Broadcast address:3.1.7.255
Netmask:255.255.255.0
Maximum Transfer Unit Size:1500
Metric:1
Packets Received: 1875
Input Errors: 0
Input Packets Dropped: 0
Input Packets Overruns: 0
Input Packets Frames: 0
Packet Sent: 5221
Output Errors: 0
Output Packets Dropped: 0
Output Packets Overruns: 0
Output Packets Carrier: 0
Output Queue Length:0
Collisions: 0
Flags:UP BROADCAST RUNNING MASTER MULTICAST

SE# show interface portChannel 2
Interface PortChannel 2 (2 physical interface(s)):
GigabitEthernet 1/0 (active)
GigabitEthernet 2/0 (active)
--------------------
Type:Ethernet
Ethernet address:00:30:48:33:01:26
Internet address:3.1.8.200
Broadcast address:3.1.8.255
Netmask:255.255.255.0
Maximum Transfer Unit Size:1500
Metric:1
Packets Received: 21
Input Errors: 0
Input Packets Dropped: 0
Input Packets Overruns: 0
Input Packets Frames: 0
Packet Sent: 0
Output Errors: 0
Output Packets Dropped: 0
Output Packets Overruns: 0
Output Packets Carrier: 0
Output Queue Length:0
Collisions: 0
Flags:UP BROADCAST RUNNING MASTER MULTICAST

SE# show statistics wmt streamstat
Detailed Stream Statistics
==========================
Incoming Streams:
=================
Stream-Id Type Source State Bytes_Recd Duration Bandwidth Url_Requested
Outgoing Streams:
=================
Client-IP Type Transport Source State Pkts_sent Bytes_sent Duration BW Server-IP Filename
2.224.22.50 VOD RTSPT LOCAL Play 5221 7312466 42 502 local 100kbps.wmv

SE# show cms info
CDN information :
Model = SE200
Node Id = 424
```

```
Device Mode = se
Current CDSM address = 4.0.5.5
Registered with CDSM = 4.0.5.5
Status = Online
Time of last config-sync = Thu May 3 16:56:54 2007
CMS services information :
Service cms_se is running
```

In the Devices Table page on the CDSM (**Devices > Devices**), the SE or SR status should be "Online." The IP address for the device always shows the IP address of the primary interface.

Configuring Last-Resort Routing

Last-resort routing is applicable when load-based routing is enabled and all Service Engines have exceeded their thresholds or all Service Engines in the domain are offline. The Service Router can redirect requests to a configurable alternate domain when all Service Engines serving a client network region are overloaded.

Note If the last-resort domain is not configured and the Service Engine thresholds are exceeded, requests are redirected to the origin server.

To configure last-resort routing use the **service-router** global configuration command, where *domain* is the service routing domain name, and *alternate* is where to route requests.

service-router lastresort domain *domain* **alternate** *alternate*

In the example below, srfqdn.cisco.com is the service routing domain name, and www.cisco.com is the alternate domain name.

```
SR(config)# service-router ?
    lastresort Configure lastresort domain
    leastloaded Enable Load Based Routing
    location-based-routing Configure location based routing
SR(config)# service-router lastresort ?
    domain Configure domain
SR(config)# service-router lastresort domain srfqdn.cisco.com ?
    alternate Configure alternate domain
SR(config)# service-router lastresort domain srfqdn.cisco.com alternate ?
    WORD Configure alternate domain name
SR(config)# service-router lastresort domain srfqdn.cisco.com alternate www.cisco.com ?
    <cr>
SE(config)# service-router lastresort domain srfqdn.cisco.com alternate www.cisco.com
```

Configuring Standby Interfaces

You can configure one ore more interfaces to act as a backup interface (a standby interface) for another interface on a Service Engine. This feature is called standby interface support. Standby groups, which are logical groups of interfaces, are used to implement this feature. When an active network interface fails (because of cable trouble, Layer 2 switch failure, high error count, or other failures) and that interface is part of a standby group, a standby interface can become active and take the load off the failed interface.

A standby group must have at least two interfaces. Interfaces that are part of a standby group are called member interfaces. After you create a standby group, you define which interfaces should be assigned to this logical group. As part of defining the member interfaces, you specify the priority of each member interface in a standby group. The member interface with the highest assigned priority is the active interface for that particular standby group. If the active interface fails, the operational member interface with the next highest priority in the standby group comes up, and so forth. If all member interfaces of a particular standby group are down and then one of the member interfaces comes up, the CDS software detects this situation and brings up the standby group on the member interface that just came up.

The failure or failover of member interfaces within a standby group triggers alarms and traps (if alarms and traps are enabled on the Service Engine). Alarms are sent out when failover occurs between member interfaces in a standby group. Specifically, minor alarms are sent out when member interfaces fail, and these alarms are cleared automatically when the interface failover has been successfully completed. Major alarms are sent out if the standby group goes down (no member interface in a standby group can be brought up).

Note A physical interface can belong to more than one standby group, and a single interface can act as a standby interface for more than one standby group.

To configure standby interfaces, interfaces are logically assigned to standby groups. The following rules define the standby group relationships:

- Each standby group is assigned a unique standby IP address, shared by all member interfaces of the standby group. The IP address of the standby group is shared among the member interfaces; however, only the active interface of the standby group uses this shared IP address at any one time. This shared IP address is configured as an alias on the active interface.
- The duplex and speed settings of the member interfaces can be configured for better reliability.
- If a physical interface is a member of a port-channel group, it cannot join a standby group. If a physical interface is a member of a standby group, it cannot join a port-channel group.
- The maximum number of standby groups on a Service Engine is four.

Note Interface IP addresses and standby group IP addresses must be on different subnets to ensure reliable operation. You can use dummy IP addresses in the private address space to serve as interface primary IP addresses, and use the real Service Engine IP address to serve as the standby group IP address in a different subnet to satisfy this requirement. When dummy IP addresses are used, these interface IP addresses serve only as substitutes to bring up the interface. For example, the Service Engine interface requires an IP address on an interface for initialization. Make sure to configure the interface default gateway using the **ip default-gateway** global configuration command instead of the **ip route** command.

- Each interface in a standby group is assigned a priority. The operational interface with the highest priority in a standby group is the active interface. Only the active interface uses the group IP address.
- The priority of an interface in a standby group can be changed at run time. The member interface that has the highest priority after this change becomes the new active interface (the default action is to preempt the currently active interface if an interface with higher priority exists).
- The maximum number of errors allowed on the active interface before the interface is shut down and the standby is brought up is configured with the **errors** option, which is disabled by default.

Tip If an interface belongs to more than one standby group, you can configure the interface with a different priority in each standby group for better load balancing. For example, interfaces FE 0/0 and FE 0/1 are both in standby group 1 and in standby group 2. If you configure FE 0/0 with the highest priority in standby group 1 and configure FE 0/1 with the highest priority in standby group 2, standby group 1 will use FE 0/0 as the active interface, while standby group 2 will use FE 0/1 as the active interface. This configuration allows each interface to back up the other one, if one of them fails.

Use the **interface standby** global configuration command to create standby groups on Service Engines.

Note Unlike port channels, standby groups do not support IP ACLs at a group level. However, you can configure a member interface of a standby group to support an IP ACL at the interface level. For example, you can individually configure the two member interfaces of Standby Group 1 (the Fast Ethernet slot 0/port 0 interface and the Fast Ethernet slot 0/port 1 interface) to support an IP ACL named ACL1 but you cannot configure the Standby Group 1 to support ACL1.

To configure an interface to be a backup for another interface, use the **standby** interface configuration command. To restore the default configuration of the interface, use the **no** form of this command.

standby *group_number* {**description** *text* | **errors** *max-errors* | **ip** *ip-address netmask* | **priority** *piority_level* | **shutdown**}

no standby *group_number* {**description** *text* | **errors** *max-errors* | **ip** *ip-address netmask* | **priority** *priority_level* | **shutdown**}

Syntax Description

group_number	Standby group number (1–4).
description	(Optional) Sets the description for the specified interface.
text	Description for the specified interface. The maximum length of the description text is 240 characters.
errors	Sets the maximum number of errors allowed on the active interface before the interface is shut down and the standby interface is brought up. This option is disabled by default.
max-errors	Maximum number of errors (o-4294967295).
ip	Sets the IP address for the specified standby group (Standby Group 1, 2, 3, or 4).
ip-address	IP address of the specified standby group (Standby Group 1, 2, 3, or 4). The group IP address and netmask of a standby group must be configured on all of the member interfaces.
netmask	Netmask of the specified standby group (Standby Group 1, 2, 3, or 4).
priority	Sets the priority of the member interface within a standby group. The priority of a member interface can be changed at run time. The member interface that has the highest priority after this change becomes the new active interface (the default action is to preempt the currently active interface if an interface with higher priority exists).

priority_level	Each member interface is assigned a priority number. The member interface with the highest priority number is the active interface for that standby group. Only the active interface uses the group IP address. If the **priority** option is specified without a priority number, the default value of 100 is used.
shutdown	(Optional) Shuts down the specified standby group (Standby Group 1, 2, 3, or 4). You can shut down a standby group even if you have not configured a group IP address of the standby group. **Note** When a standby group is shut down, all of the alarms previously raised by this standby group are cleared.

Examples

The following example configures three Fast Ethernet interfaces to be part of the same standby group, with interface 3/0 as the active interface:

```
Console(config-if)# interface fastEthernet 3/0 standby 1 ip 172.16.10.10 255.255.254.0
Console(config-if)# interface fastEthernet 3/1 standby 1 ip 172.16.10.10 255.255.254.0
Console(config-if)# interface fastEthernet 3/2 standby 1 ip 172.16.10.10 255.255.254.0
Console(config-if)# interface fastEthernet 3/0 standby 1 priority 300
Console(config-if)# interface fastEthernet 3/1 standby 1 priority 200
Console(config-if)# interface fastEthernet 3/2 standby 1 priority 100
Console(config-if)# interface fastEthernet 3/0 standby 1 errors 10000
Console(config-if)# interface fastEthernet 3/1 standby 1 errors 10000
Console(config-if)# interface fastEthernet 3/2 standby 1 errors 10000
```

The following example displays information about the standby group configuration by entering the **show standby** EXEC command. In the following sample command output, one standby group (Standby Group 1) is configured on this Service Engine. The command output also shows which member interface is the active interface. In this case, the active interface is the Fast Ethernet slot 3/port 0 interface.

```
ServiceEngine# show standby
Standby Group:1
IP address: 172.16.10.10, netmask: 255.255.254.0
Maximum errors allowed on the active interface: 10000
       Member interfaces:
              FastEthernet 3/0 priority: 300
              FastEthernet 3/1 priority: 200
              FastEthernet 3/2 priority: 100

       Active interface: FastEthernet 3/0
```

Note To display information about a specific standby group configuration, enter the **show interface standby** *group_number* EXEC command.

The following example creates a standby group, Standby Group 1:

```
ServiceEngine# configure
ServiceEngine(config)# interface standby 1
ServiceEngine(config-if)#
```

The following example assigns a group IP address of 10.10.10.10 and a netmask of 255.0.0.0 to Standby Group 1.

```
ServiceEngine(config-if)# ip address 10.10.10.10 255.0.0.0
ServiceEngine(config-if)# errors 500
```

The following example shows how to add two Fast Ethernet interfaces to Standby Group 1 and then assign each of these member interfaces a priority within the group:

1. Add a Fast Ethernet interface (slot 0/port 0) to Standby Group 1 and assign a priority of 150.

```
ServiceEngine(config)# interface FastEthernet 0/0
ServiceEngine(config-if)# standby 1 priority 150
```

2. Add a second Fast Ethernet interface (slot 0/port 1) to Standby Group 1 and assign a priority of 100.

(the default value):

```
ServiceEngine(config)# interface FastEthernet 0/1
ServiceEngine(config-if)# standby 1
ServiceEngine(config-if)# exit
ServiceEngine(config)#
```

Because Fast Ethernet 0/0 is assigned the highest priority (a priority number of 150) of all the member interfaces in the group, it will be chosen as the active interface for the group if it can be brought up.

The following example removes the Fast Ethernet slot 0/port 1 interface) from Standby Group 1 using the **no** form of the **standby** command:

```
ServiceEngine(config)# interface FastEthernet 0/1
ServiceEngine(config-if)# no standby 1
ServiceEngine(config-if)# exit
ServiceEngine(config)#
```

The following example shows how to shut down Standby Group 1. When a standby group is shut down, all of the alarms previously raised by this standby group are cleared.

```
ServiceEngine(config)# interface standby 1
ServiceEngine(config-if)# shutdown
ServiceEngine(config)# exit
```

The following example shows how to tear down Standby Group 1:

```
ServiceEngine(config)# interface standby 1
ServiceEngine(config-if)# no ip address 10.10.10.10 255.0.0.0
Please remove member interface(s) from this standby group first.
ServiceEngine(config)# interface GigabitEthernet 2/0
ServiceEngine(config-if)# no standby 1
ServiceEngine(config-if)# exit
ServiceEngine(config)# interface standby 1
ServiceEngine(config-if)# no ip address 10.10.10.10 255.0.0.0
ServiceEngine(config-if)# exit
ServiceEngine(config)# no interface standby 1
ServiceEngine(config)# exit
```

Other CLI Commands

This section lists other CLI commands that may be useful.

CDNFS cleanup Command

To manage the CDS network file system (cdnfs), use the **cdnfs** EXEC command.

cdnfs {**browse** | **cleanup** {**info** | **start** | **stop**} | **delete-unused-ecdnfs-files**}

There are no default behavior or values.

Syntax Description

browse	Browses the cdnfs directories and files.
cleanup	Cleans up the unwanted entries in the cdnfs.
info	Summarizes the information about unwanted entries without starting the cleanup process.
start	Starts the cleanup of unwanted entries in the cdnfs.
stop	Stops the cleanup of unwanted entries in the cdnfs.

The CDNFS cleanup command is an EXEC command.

The CDS network file systems (cdnfs) stores the prefetched CDS network content by all supported protocols. The **cdnfs cleanup** command cleans up the content of deleted channels from the acquisition distribution database. In certain cases, the Content Acquirer is not notified by the Centralized System (CMS) about deleted delivery services, and it fails to clear all unified name space (UNS) cases, the **cdnfs cleanup** EXEC command can be used to clean up all UNS content associated with deleted delivery services.

Note You can use **cdnfs cleanup start** to clean up the orphan content. Orphan content is content that is not associated with any delivery service to which a Service Engine is subscribed.

The **cdnfs browse** command is an interactive command and has the following interactive commands used to view the CDS network files and directories:

```
SE# cdnfs browse
------ CDNFS interactive browsing ------
dir, ls: list directory contents
cd,chdir: change current working directory
info: display attributes of a file
more: page through a file
cat: display a file
exit,quit: quit CDNFS browse shell
```

The **cdnfs cleanup** command synchronizes the state of the acquisition and distribution database with the content stored on the cdnfs. You should use this command after replacing a failed disk drive.

Disk Commands

The **disk** commands in the EXEC mode allow you to perform disk configuration and maintenance commands.

- **disk mark** *diskname* {**bad** | **good**}
- **disk recover-system-volumes**
- **disk reformat** *diskname*
- **disk unuse** *diskname*

Syntax Description

mark	Marks a disk drive as good or bad.
diskname	Name of the disk to be added (disk01, disk02, and so on.)
recover-system-volumes	Recovers the system disk.

reformat	Performs a low-level format of the SCSI, IDE, or SATA disks and remaps disk errors.
diskname	Name of the disk to be added (disk01, disk02, and so on.)
unuse	Stops applications from using a disk drive.
diskname	Name of the disk to be added (disk01, disk02, and so on.)

The **disk** commands are EXEC commands.

The **disk reformat** *diskname* EXEC command performs a low-level format of the SCSI, IDE, or SATA disks. This command erases all of the content on the disk.

If a disk drive continues to report a failure after you have used the **disk reformat** command, you must replace the disk drive.

Caution Be careful when using the **disk reformat** *diskname* command because this command causes all content on the specified disk to be deleted.

The **disk unuse** command stops and restarts all applications that are currently using the specified disk drive (for example, disk02 [/local/local2] or disk03), unmounts, and deletes all the partitions on the specified disk. Using this command unmounts all file systems.

Note For information about replacing a disk, see the "Disk Maintenance" section on page 8-26.

cache content Command

To set the number of cache entries in the CDS network file system (cdnfs), use the **cache content** global configuration command.

cache content max-cached-entries *1-10000000*

Syntax Description

content	Browses the cdnfs directories and files.
max-cached-entries	Cleans up the unwanted entries in the cdnfs.
1-10000000	Summarizes the information about unwanted entries without starting the cleanup process.

Maximum cache entries is set to 3000000 by default.

The CDS, by default, allows a maximum of three million cached entries, regardless of the amount of space available in CDNFS. Use this command to restrict the maximum number of contents in the CDS.

APPENDIX F

URL Signing and Validation

This appendix describes the URL signing and validation method for the Cisco Internet Streamer CDS. This appendix contains the following sections:

- Introduction, page F-1
- Configuring the CDS for URL Signing, page F-2
- URL Signing Script, page F-2
- Understanding the Signing Procedure , page F-5

Introduction

The Cisco Internet Streamer CDS accepts and fulfills requests for video content from client devices in the form of content URLs. Content and service providers, in order to protect their copyright and fulfill their licensing obligations, often need to restrict access to content and limit viewing times. Basic authentication and authorization at the portal (for example, username and passwords) can help achieve this objective by restricting content access to authorized users. However, because URLs are inherently open, users (once authenticated at the portal) could potentially share these content URLs with other possibly unauthorized users, or continue to access the content beyond the allotted time.

Cisco Internet Streamer CDS 2.2 provides the infrastructure to sign and validate content URLs, restricting access to some users and limiting viewing times.

URL Signing Components

One of the easiest ways to restrict content access to a particular user is to embed, within the content URL, the client IP address of the user for whom the content access was authorized. Similarly, to ensure that the content expires after a predetermined time, an expiry timestamp could be embedded. These values can then be validated against the actual client sending the request and the current time at the Service Engine serving the request. If either of the two validations fail, the request is rejected.

However, because any of these strings in the URL could potentially be edited manually and circumvented by any knowledgeable user, it is important to generate and attach a signature to the URL. This can be achieved by attaching a keyed hash to the URL, using a secret key shared only between the signer (the portal) and the validating component (CDS).

CDS has incorporated an open and well-documented signing mechanism that uses standard hashing schemes. The URL signing mechanism offers the flexibility to either use the provided signing script, or you can develop a signing application in the platform or language of your choice, as long as it adheres to the specified format.

For signing and validation of the URL, the CDS relies on a set of one or more secret keys shared between the portal and the devices within the CDS.

Supported Protocols and Media

The URL signing and validation is supported across all CDS protocol engines; Windows Media Engine, Movie Streamer Engine, Flash Media Streaming Engine, and Web Engine.

Configuring the CDS for URL Signing

To enable validation of URLs in the CDS, the following tasks must be completed on all participating Service Engines:

- Configure shared secret keys
- Configure pattern-lists to match URLs, domain names , or both
- Configure rules to validate URLs matching the above pattern-lists
- Enable rules processing

Details on these configurations are available in the "Configuring URL Signing" section on page 4-23 and the "Configuring Service Rules" section on page 4-15.

The CDS URL signing infrastructure supports multiple keys. Different pieces of content, with different URLs, can be signed by different keys. Keys are stored as a key matrix and identified (indexed) by a key ID owner and a key ID number.

URL Signing Script

At the portal, URLs can be signed for a particular user (client IP address) and expiry time using a URL signing script. The URL signing script example included in this appendix requires Python 2.3.4 or higher.

Example of a Python URL Signing Script

The following simple Python script demonstrates how to construct and sign URLs for use with the Internet Streamer CDS Release 2.2. This example script produces signatures compliant with the format used by the Internet Streamer CDS.

Depending on where the python binary is installed, you may need to modify the first line of the script. The first line is only necessary if you plan to run the script as an executable. However, if you run the script using the python interpreter, as documented in the "Running a Python URL Signing Script" section on page F-3, the first line is not required.

```
#!/usr/local/bin/python
import md5
import socket
```

```
import time
import sys

def sign_url(url,key):
    """ Signs url using key and returns the signed URL with the signature appended. """
    # Generate a MD5 hash of the key string (not the url)
    foo = md5.new(key)
    # Update the hash generated with the url string
    # This effectively means concatenating key and url and generating
    # a hash for the two
    foo.update(url)
    # Get the digest in hex format (human readable)
    return url+foo.hexdigest()

def usage():
    """ Prints usage for the URL signing script. """
    print "Usage:"
    print "python cds-ims-urlsign.py <url> <client-ip> <expiry-delay-seconds>
    <key-id-owner> <key-id-number> <key>"
    print "Example:"
    print "python cds-ims-urlsign.py rtsp://abc.com/content/Apocalypto.mov 171.71.50.123
    120 1 2 pl023MDQ1k"

if __name__ == "__main__":
    """ Prints signed URL """

    if len(sys.argv) < 7:
        usage()
        sys.exit(2)
    url = sys.argv[1] # URL
    client_ip = sys.argv[2]
    delay_seconds = sys.argv[3] # Number of seconds after which URL expires
    ko = sys.argv[4] # Key ID Owner
    kn = sys.argv[5] # Key ID Number
    key = sys.argv[6] # Key

    # Set expiry time as current time (seconds since epoch) + delay
    et = time.time() + int(delay_seconds)
    expires = str(int(et))


    if url.find('?')==-1:
        url2 = url + "?"
    else:
        url2 = url + "&"

# This string format is fixed and should not be modified.
# Note that we sign even the "&US=" part that will point to the signature
url2 = url2 + "IS=0"+"&ET="+expires+"&CIP="+client_ip+"&KO="+ko+"&KN="+kn+"&US=";

url3 = sign_url(url2,key)

#print url
#print url2
print url3
```

Running a Python URL Signing Script

The example script, call it “cds-ims-urlsign.py,” can be used as follows:

python cds-ims-urlsign.py *url client-ip expiry-delay-seconds key-id-owner key-id-number key*

url	URL to sign.
client-ip	IP address of the client for which this URL is being signed, in dotted decimal format (A.B.C.D). The signed URL will be rejected if sent from any other client if signature validation is enabled.
expiry-delay-seconds	Seconds (from now) when the URL expires. The request will be rejected if the time period has passed when the URL is validated at the device. See the "Importance of Device Synchronization" section on page F-4.
key-id-owner	This argument provides the first index into the key matrix.
key-id-number	This argument provides the second index into the key matrix.
key	Shared secret key corresponding to this ordered pair (key-id-owner, key-id-number).

In addition to the above six variables, the current time is used to generate the URL signing, so even if the same values were used for the above six variables, the signed URL would be different.

To use the URL signing script on the URL "rtsp://cisco.com/content/CiscoCDS.mov," for the client IP address of 171.71.50.123, with expiry delay of 120 seconds, key-id owner of 1, key-id-number of 2, and a key of kwnx90KGP, enter the following:.

python cds-ims-urlsign.py rtsp://cisco.com/content/CiscoCDS.mov 171.71.50.123 120 1 2 kwnx90KGP

The signed URL is the following:

```
rtsp://cisco.com/content/CiscoCDS.mov?IS=0&ET=1209422976&CIP=171.71.50.123&KO=1&KN=2&US=f0
8b56f46075813e44b2d4888628a471
```

Note The above signed URL is only an example. The MD5 algorithm generates a different message digest each time. For more information on the MD5 algorithm see the IETF RFC 1321.

Importance of Device Synchronization

URL expiry time validation relies on the assumption that the clocks are synchronized on the server running the signing application and the Service Engines validating the URL. Use of Network Time Protocol (NTP) on all devices, including the device running the signing application or script, is highly recommended.

It is not sufficient to merely have the same local times on two devices while their time zones differ.

For example, the following two devices are not synchronized:

- Device 1:
  - Local Time: 11:00:59 PM, October 12, 2008
  - Time Zone: PST
- Device 2:
  - Local Time: 11:00:59 PM, October 12, 2008
  - Time Zone: EST

The following two devices are synchronized:

- Device 1:
  - Local Time: 11:00:59 AM, October 12, 2008
  - Time Zone: PST
- Device 2:
  - Local Time: 2:00:59 PM, October 12, 2008
  - Time Zone: EST

Understanding the Signing Procedure

To customize the URL signing script for your portal, or to write your own signing application in the platform and language of your choice, and still be able to validate URLs within the CDS, follow the steps explained in this section.

The URL signing script performs these steps when processing an unsigned URL:

1. Checks if the URL already contains a query string.

   If the URL does not contain a query string, appends a question mark (?).

   If the URL does contain a query string, appends an ampersand (&).
2. Appends the string **IS=0**. This string is for legacy support with some CDS components that use both internal (within CDS) and external (**portal**) signing mechanisms.
3. Appends the string **&ET=**.
4. Gets the current time in seconds since epoch (as an integer). Adds the expiry time in seconds as an integer and appends this integer.
5. Appends the string **&CIP=**.
6. Appends the requesting client IP address, using dotted decimal format.
7. Appends the string **&KO=**.
8. Appends the key ID owner corresponding to the key being used.
9. Appends the string **&KN=**.
10. Appends the key ID number corresponding to the key being used.
11. Appends the string **&US=**.
12. Stores this as url2; for example:

    “rtsp://cisco.com/content/CiscoCDS.mov?IS=0&ET=1209422976&CIP=171.71.50.123&KO=1&KN=2&US=”
13. Generates an MD5 hash of the key being used.
14. Updates the generated hash with url2.
15. Converts the hash to its equivalent human readable hex digest; for example:

    f08b56f46075813e44b2d4888628a471
16. Appends the hex digest to url2. The URL signing is complete.

APPENDIX G

Testing the Internet Streamer CDS

This appendix covers the steps to test the CDS by using the different media players. This appendix covers the following topics:

- Testing the Web Engine, page G-1
- Testing the Windows Media Streaming Engine, page G-7
- Testing the Movie Streamer Engine, page G-11
- Testing the Flash Media Streaming Engine, page G-19

The CDS network topology example used in these procedures consists of the following devices:

- 2 CDE200s configured as Service Engines (SEs)
  - NE-DEMO-SE1 — Tier 1 location
  - NE-DEMO-SE2 — Tier 2 location
- 1 CDE200 configured as a Service Router (SR)
  - NE-DEMO-SR — Tier 2 location
- 1 CDE100 configured as a Content Delivery System Manager (CDSM)
  - NE-DEMO-CDSM

Testing the Web Engine

This section consists of the following procedures:

- Testing Preingested Web Content
- Testing Dynamically Ingested Web Content

Testing Preingested Web Content

Note Content must be preingested. See the "Configuring Delivery Services" section on page 5-1 for more information.

Step 1 In a web browser on a client PC, enter the URL of the preingested HTML content using the Service Routing Domain Name in the URL. In the example, this is "http://rfqdn.cds.com/test_prepos/test.html." Client requests are directed to this domain name and are then redirected by the Service Router to the content on the Service Engine. Figure G-1 shows the initial URL on the left and the redirection on the right.

Figure G-1 URL Redirection for Preingested Content

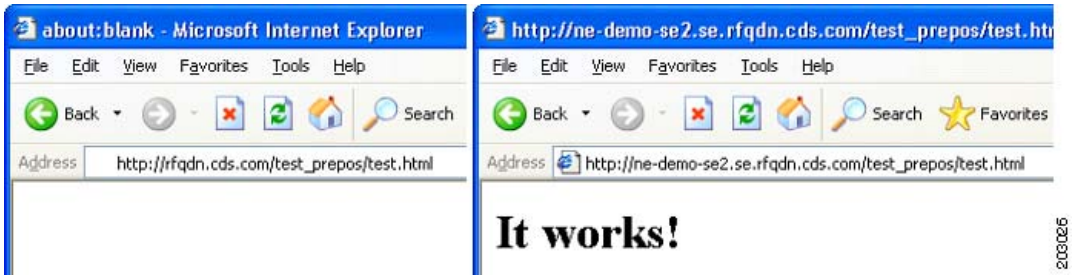

Step 2 On the Service Engines, verify that the request was served as a preingested hit. View the HTTP request statistics by starting an SSH session and entering the **show statistics http requests** command. In this case, the Service Router redirected the request to NE-DEMO-SE2, which served the request.

```
NE-DEMO-SE1# show statistics http requests

    Caching Statistics
    ------------------
            Cache Hits : 0
            Cache Miss : 0
         Range Requests : 0
Partial HIts - Live fill :0
    Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
        Preposition Hits : 0
             Reply Meta :  0
         Alternate Media : 0

    Others
    ------
            Num Lookups : 0
           Lookup Errors :0
Streaming redirected requests : 0
       WMT Liveness requests :0
Hierarchical Cache Liveness requests : 0

                   Total       % of requests
                   -----       -------------
              Client Errors :    0    0.00
              Server Errors :    0    0.00
          Http/0.9 Requests :    0    0.00
          Http/1.0 Requests :    0    0.00
          Http/1.1 Requests :    0    0.00
     Http Invalid Requests :     0    0.00
                   Blocked :     0    0.00
                   Allowed :     0    0.00
```

```
NE-DEMO-SE2# show statistics http request

Caching Statistics
------------------
           Cache Hits : 0
           Cache Miss : 0
       Range Requests : 0
Partial HIts - Live fill :0
    Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
        Preposition Hits : 0
            Reply Meta :  0
         Alternate Media : 0

    Others
    ------
           Num Lookups : 1
          Lookup Errors :0
Streaming redirected requests : 0
      WMT Liveness requests :0
Hierarchical Cache Liveness requests : 0

                  Total       % of requests
                  -----       -------------
            Client Errors :     0    0.00
            Server Errors :     0    0.00
        Http/0.9 Requests :     0    0.00
        Http/1.0 Requests :     0    0.00
        Http/1.1 Requests :     0  100.00
    Http Invalid Requests :     0    0.00
                  Blocked :     0    0.00
                  Allowed :     0  100.00
```

Testing Dynamically Ingested Web Content

Step 1 In a web browser on the client PC, enter the URL of non-preingested HTML content on the Service Router. This is content that exists on the origin server or some other server that is accessible but not yet preingested. In the example, the origin server has a directory “test_cache” with a content object “test.html.” Figure G-2 shows the initial URL on the left and the redirection on the right.

Figure G-2 URL Redirection for Non-Preingested Content

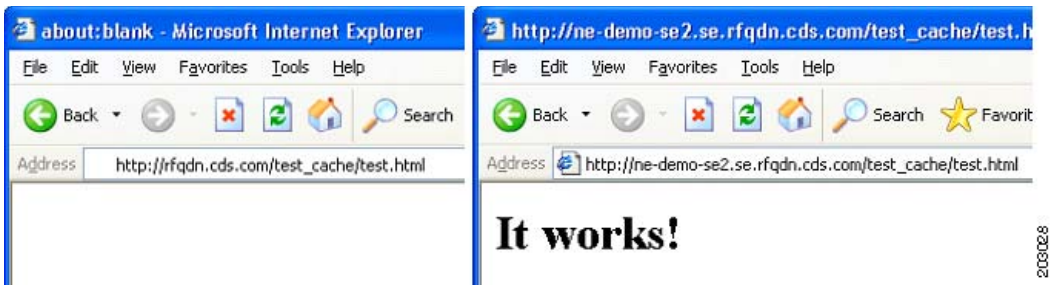

This is a cache miss scenario. Neither Service Engine had the content preingested, so the content is acquired by NE-DEMO-SE1 (the Content Acquirer). The content is then cached and replicated to NE-DEMO-SE2 (the receiver, which also happens to be the SE that is serving this client request). NE-DEMO-SE2 then serves the request (as visible by the new URL in Figure G-2), having cached the content as well.

Step 2 View the HTTP request statistics by entering the **show statistics http requests** command.

```
NE-DEMO-SE1# show statistics http requests

Caching Statistics
------------------
           Cache Hits : 0
           Cache Miss : 1
        Range Requests : 0
Partial HIts - Live fill :0
    Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
    Preposition Hits : 0
         Reply Meta :  0
     Alternate Media : 0
        Others
        ------
    Num Lookups : 2
Lookup Errors :  0
Streaming redirected requests : 0
WMT Liveness requests :0
Hierarchical Cache Liveness requests : 4

                 Total    % of requests
                 -----    --------------
 Client Errors :     0     0.00
 Server Errors :     0     0.00
Http/0.9 Requests :  0     0.00
Http/1.0 Requests :  0     0.00
Http/1.1 Requests :  5   100.00
Http Invalid Requests :0   0.00
        Blocked :    0     0.00
        Allowed :    5   100.00

NE-DEMO-SE2# show statistics http requests

Caching Statistics
------------------
    Cache Hits : 0
    Cache Miss : 1
Range Requests : 0
Partial HIts - Live fill :0
Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
    Preposition Hits :1
         Reply Meta : 0
     Alternate Media :0

        Others
        ------
    Num Lookups :3
```

```
Lookup Errors :0
Streaming redirected requests : 0
WMT Liveness requests :0
Hierarchical Cache Liveness requests : 0

                Total   % of requests
                -----   -------------
Client Errors :     0     0.00
Server Errors :     0     0.00
Http/0.9 Requests : 0     0.00
Http/1.0 Requests : 0     0.00
Http/1.1 Requests : 2   100.00
Http Invalid Requests :0  0.00
        Blocked :   0     0.00
        Allowed :   2   100.00
```

Step 3 Verify that the content was cached properly on NE-DEMO-SE1 and NE-DEMO-SE2 by entering the **show cache content** command.

```
NE-DEMO-SE1# show cache content
Max-cached-entries is set as 3000000
Number of cal cached assets: 1
-----------------------------------------------
Priority   Size    URL
-----------------------------------------------
5.41532e-01 44    http://ofqdn.cds.com/test_cache/test.html

NE-DEMO-SE2#show cache content
Max-cached-entries is set as 3000000
Number of cal cached assets: 1
-----------------------------------------------
Priority    Size   URL
-----------------------------------------------
6.20339e-01 44     http://ofqdn.cds.com/test_cache/test.html
```

Step 4 Clear the cache content in the web browser in order to make sure future requests are handled by the Service Engines instead of from the browser's local cache.

Step 5 In the web browser, request the same content again (http://rfqdn.cds.com/test_cache/test.html). Figure G-3 shows the initial URL request on the left and the new URL on the right.

Figure G-3 Show Statistic Requests on NE-DEMO-SE1—Cached Content

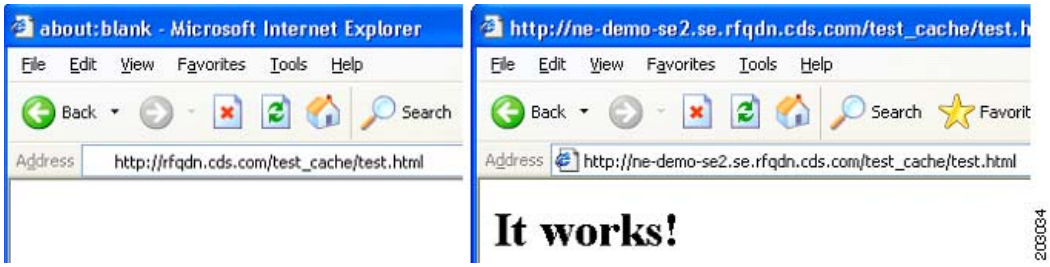

Step 6 View the HTTP request statistics again by entering the **show statistics http requests** command.

```
NE-DEMO-SE1# show statistics http requests

Caching Statistics
------------------
    Cache Hits : 0
    Cache Miss : 1
Range Requests : 0
Partial HIts - Live fill :0
```

```
Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
    Preposition Hits : 0
         Reply Meta :  0
     Alternate Media : 0
        Others
        ------
    Num Lookups : 2
Lookup Errors :  0
Streaming redirected requests : 0
WMT Liveness requests :0
Hierarchical Cache Liveness requests : 18

                Total    % of requests
                -----    --------------
 Client Errors :     0     0.00
 Server Errors :     0     0.00
Http/0.9 Requests :  0     0.00
Http/1.0 Requests :  0     0.00
Http/1.1 Requests :  19  100.00
Http Invalid Requests :0   0.00
        Blocked :    0     0.00
        Allowed :    19  100.00

NE-DEMO-SE2# show statistics http requests


Caching Statistics
------------------
    Cache Hits : 1
    Cache Miss : 1
Range Requests : 0
Partial HIts - Live fill :0
Partial HIt - Refill : 0
Partial Caching - Bypassed : 0

Prepositioned Content Statistics
--------------------------------
    Preposition Hits : 1
         Reply Meta :  0
     Alternate Media : 0
        Others
        ------
    Num Lookups : 5
Lookup Errors :  0
Streaming redirected requests : 0
WMT Liveness requests :0
Hierarchical Cache Liveness requests : 0

                Total    % of requests
                -----    --------------
 Client Errors :     0     0.00
 Server Errors :     0     0.00
Http/0.9 Requests :  0     0.00
Http/1.0 Requests :  0     0.00
Http/1.1 Requests :  3   100.00
Http Invalid Requests :0   0.00
        Blocked :    0     0.00
        Allowed :    3   100.00
```

In this case, NE-DEMO-SE2 served the request, and it is a cache hit scenario. The content was cached from the previous attempt, and now the same content is served from cache.

Testing the Windows Media Streaming Engine

This section consists of the following procedures:

- Testing Preingested Windows Media Content
- Testing Dynamically Ingested Windows Media Content
- Testing Windows Media Live Content Playback

Testing Preingested Windows Media Content

Note Content must be preingested. See the "Configuring Delivery Services" section on page 5-1 for more information.

Step 1 On the client PC, start the Windows Media Player program.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL of the preingested content in the Open URL dialog box (Figure G-4) and click **OK**. In the example, this is "rtsp://rfqdn.cds.com/test_prepos/test.wmv."

Figure G-4 Open URL Dialog Box

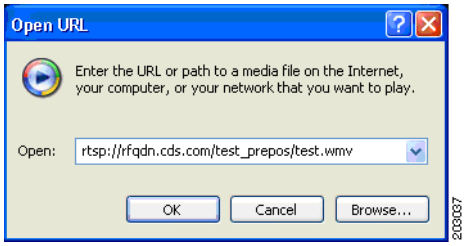

The video begins to play.

Step 4 To view the statistics on the video file, choose **View > Statistics** and then click the **Advanced** tab (Figure G-5).

Figure G-5 Windows Media Player Statistics

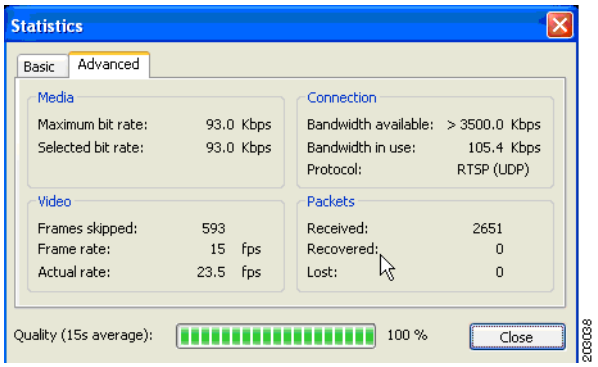

Step 5 To view the request flow, enter the **show statistics wmt streamstat** command on the SEs. In this case, the request is served from NE-DEMO-SE2.

```
NE-DEMO-SE1# show statistics wmt streamstat
Detailed Stream Statistics
====================================

Incoming Streams:
=================
Stream-Id  Type  Source State  Bytes-Recd   Duration   Bandwidth  Server-IP  Url-Requested

Outgoing Streams:
=================
Client-IP  Type  Transport Source  State Pkts-set  Bytes-sent  Duration  BW Filename Stream-id

NE-DEMO-SE2# show statistics wmt streamstat
Detailed Stream Statistics
====================================

Incoming Streams:
=================
Stream-Id  Type  Source State  Bytes-Recd   Duration   Bandwidth  Server-IP  Url-Requested

Outgoing Streams:
=================
Client-IP     Type  Transport Source  State Pkts-set  Bytes-sent  Duration  BW Filename          Stream-id
171.70.222.171 VOD  NONE      LOCAL   Setup  0        0           0         0  test_prepos/test.wmv   8895
```

Testing Dynamically Ingested Windows Media Content

Step 1 On the client PC, start the Windows Media Player program.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL of content that has not been preingested in the Open URL dialog box (Figure G-6) and click **OK**. In the example, this is “rtsp://rfqdn.cds.com/test_cache/test.wmv.”

Figure G-6 Open URL Dialog Box

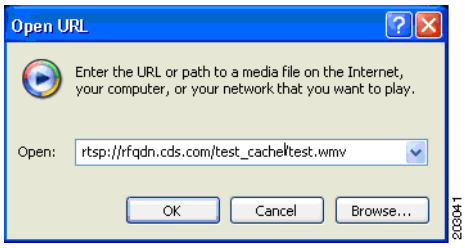

The video begins to play.

Step 4 To view the statistics on the video file, choose **View > Statistics** and then click the **Advanced** tab (Figure G-7).

Figure G-7 Windows Media Player Statistics

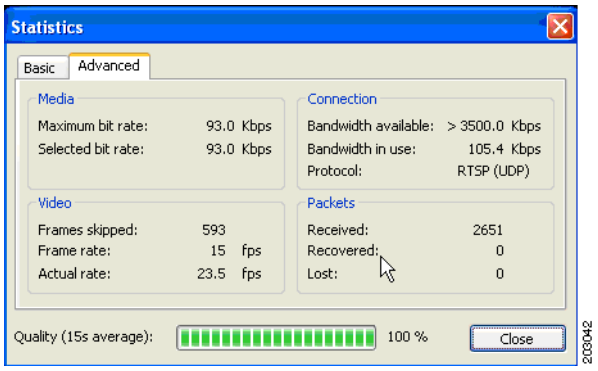

Step 5 To view the request flow, enter the **show statistics wmt streamstat** command on the SEs. In this case, the request is received from NE-DEMO-SE2.

```
NE-DEMO-SE1# clear statistics wmt
NE-DEMO-SE1# show statistics wmt streamstat
Detailed Stream Statistics
==================================

Incoming Streams:
=================
Stream-Id  Type  Source State  Bytes-Recd   Duration   Bandwidth  Server-IP  Url-Requested

Outgoing Streams:
=================
Client-IP   Type  Transport Source  State Pkts-set  Bytes-sent  Duration  BW Filename           Stream-id
3.1.4.14    VOD   RTSPT    RMT_HTTP Play  388       3104000     19        549 test_cache/test.wmv  27521

NE-DEMO-SE2# clear statistics wmt
NE-DEMO-SE2# show statistics wmt streamstat
Detailed Stream Statistics
==================================

Incoming Streams:
=================
Stream-Id  Type  Source    State  Bytes-Recd   Duration   Bandwidth  Server-IP  Url-Requested
12079      VOD   RMT_RTSP Play   2241074     23         1550       3.1.4.10
rtsp;//ofqdn.cds.com/test_cache/test.wmv
```

```
Outgoing Streams:
=================
Client-IP       Type  Transport Source   State Pkts-set  Bytes-sent  Duration  BW Filename          Stream-id
171.70.222.171  VOD   RTSPT     RMT_RTSP Play  279       2232000     13      7201 test_cache/test.wmv  12043
```

Testing Windows Media Live Content Playback

Note Each live program uses a live delivery service to deliver the live program. See the "Configuring Programs" section on page 5-26 for more information.

Step 1 On the client PC, start the Windows Media Player.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL for the live program in the Open URL dialog box (Figure G-8) and click **OK**. In the example, this is "rtsp://rfqdn.cds.com/wmtlive."

Figure G-8 Open URL Dialog Box

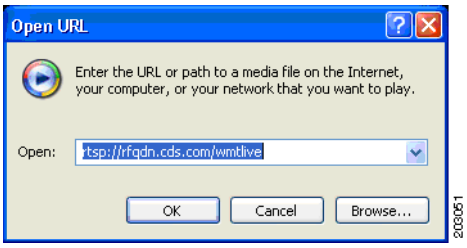

The video begins to play.

Step 4 To view the statistics on the video file, choose **View > Statistics** and then click the **Advanced** tab (Figure G-9).

Figure G-9 Windows Media Player Statistics

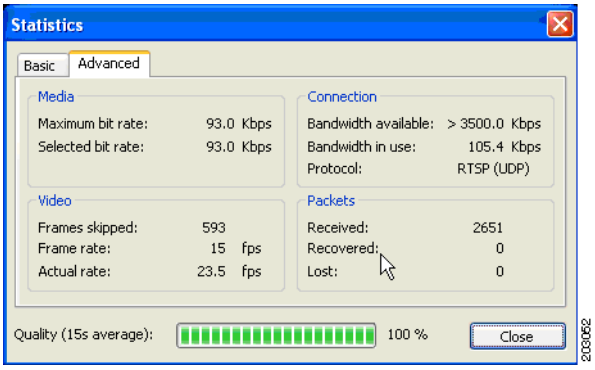

Step 5 To view the incoming and outgoing streams, enter the **show statistics wmt streamstat** command on the SEs streaming the content. In this case, the request is served from NE-DEMO-SE2.

```
NE-DEMO-SE1# show statistics wmt streamstat
Detailed Stream Statistics
====================================

Incoming Streams:
=================
Stream-Id  Type  Source    State  Bytes-Recd Duration  Bandwidth  Server-IP  Url-Requested
18872      LIVE  RMT_HTTP  Play   84150      241       288     171.70.22.171
http://171.70.222.171:0000

Outgoing Streams:
=================
Client-IP  Type  Transport Source   State Pkts-set  Bytes-sent  Duration  BW Filename  Stream-id
3.1.4.14   LIVE  RTSPT     RMT_HTTP  Play 5393      7787492     233      288 8080      18889

NE-DEMO-SE2# show statistics wmt streamstat
Detailed Stream Statistics
====================================

Incoming Streams:
=================
Stream-Id  Type  Source    State  Bytes-Recd Duration  Bandwidth  Server-IP  Url-Requested
28772      LIVE  RMT_RTSP Play   8205265   246       289         3.1.4.10
rtsp;//3.1.4.10/wmt_proxy/rtsp&ofqdn.cds.com/wmtlive/__CDS/http&171.70.111.171&8080

Outgoing Streams:
=================
Client-IP       Type  Transport Source   State Pkts-set  Bytes-sent  Duration  BW  Filename  Stream-id
171.70.222.171  LIVE  RTSPU     RMT_RTSP  Play  5582     8060408     241       288 wmtlive   28755
```

Testing the Movie Streamer Engine

Note In Release 2.0, Movie Streamer was in a demonstration state for live streaming, prefetched, cached and dynamically cached content.
In Release 2.1, live streaming is in full production. Prefetched, cached, and dynamically cached content remain in a demonstration state.
For details of live streaming performance for Movie Streamer, please refer to the Release 2.1 performance bulletin.

This section includes the following procedures:

- Preparing Movie Streamer Content for Ingest
- Testing Preingested Movie Streamer Content
- Testing Dynamically Ingested Movie Streamer Content
- Testing Movie Streamer Live Content Playback

Preparing Movie Streamer Content for Ingest

The Movie Streamer delivers hinted MPEG-4, hinted 3GPP, and hinted MOV files to clients over the Internet and mobile networks. Hinted files contain hint tracks, which store packetization information that tell the streaming server how to package content for streaming. Apple QuickTime Pro can be used to generate the hint tracks.

Step 1 Launch Apple QuickTime Pro.

Step 2 Choose **File > Open File**, and select a movie file to open.

Step 3 Choose **Windows > Show Movie Properties**. If there is a Hinted Video Track present, as shown in Figure G-10, then open the next movie file.

Figure G-10 Movie Properties Dialog Box

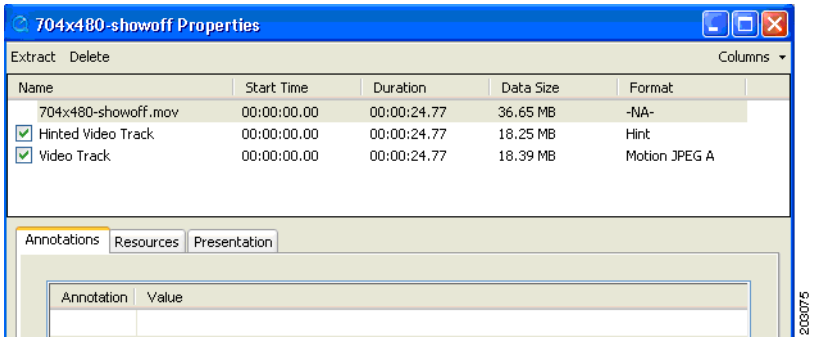

Step 4 Choose **File > Export**. The Save Exported File dialog box is displayed (Figure G-11).

Figure G-11 Save Exported File Dialog Box

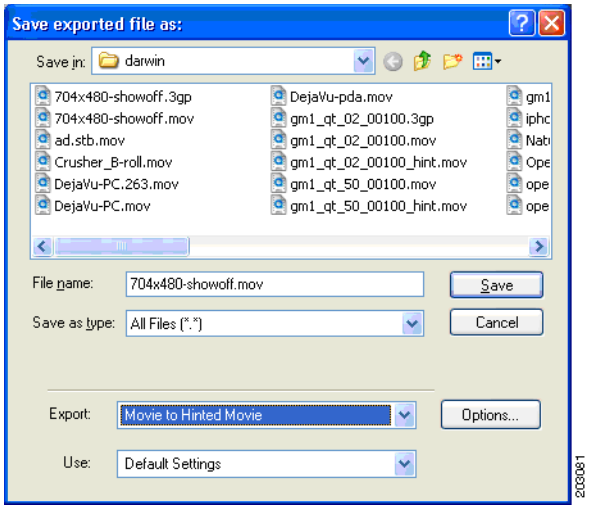

Step 5 From the **Export** drop-down list, choose **Movie to Hinted Movie**.

Step 6 Click **Options**. The Hint Exporter Settings dialog box is displayed (Figure G-12).

Figure G-12 Hint Exporter Settings Dialog Box

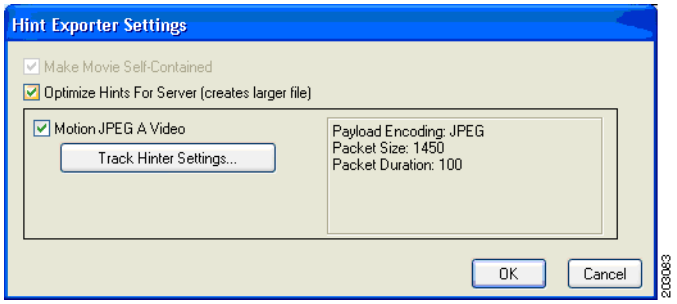

Step 7 Check the **Make Movie Self-Contained** check box and the **Optimize Hints For Server** check box.

Step 8 Click **OK**.

Step 9 Click **Save** in the Save Exported File dialog box.

The movie file is ready to be either preingested or dynamically ingested.

Testing Preingested Movie Streamer Content

Note Content must be preingested. See the "Configuring Delivery Services" section on page 5-1 for more information.

Step 1 On the client PC, start the Apple QuickTime Player.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL of the preingested content in the Open URL dialog box (Figure G-13) and click **OK**. In the example, this URL is "rtsp://rfqdn.cds.com/test_prepos/test.mov."

Figure G-13 Open URL Dialog Box

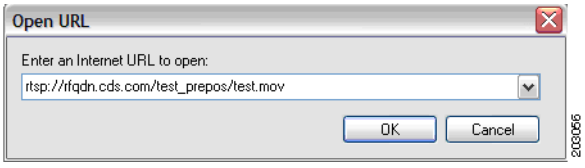

The video begins to play.

Step 4 To view the statistics on the video file, choose **Windows > Show Movie Inspector** (Figure G-14).

The current statistics are displayed and updated as the video streams and plays.

Figure G-14 QuickTime Player Statistics

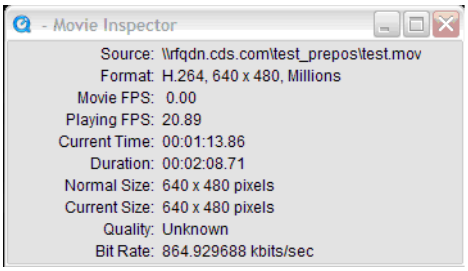

Step 5 To view the request flow, display the statistics. For CDS Releases 2.0 and 2.1, enter the **show statistics rtsp server movie-streamer all** command. For CDS Releases 2.2 and 2.3, enter the **show statistics movie-streamer all** command.

```
NE_DEMO-SE2# show statistics movie-streamer all
Movie Streamer Request Statistics
                                  Total
    ---------------------------------------
Current RTSP sessions:                    0
  Total RTSP sessions:                    0
   Current RIP connections:                   0
    Total RTP connections:                    1

CDN Related Statistics
           ----------
    Preposition Hits:                     0
         Cache Hits:                      0
         Cache Miss:                      0
       Live Requests:                     0

Cache Revalidation Statistics
                   ---------
        Fresh Content Requests:           0
          Revalidated Requests:           0

Movie Streamer Bandwidth Usage Statistics
                                  Total
---------------------------------------
Current Incoming bandwidth:               0 bps
Current Outgoing bandwidth:               0 bps
   Current Total bandwidth:               0 bps

Average Incoming bandwidth:               0 bps
Average Outgoing bandwidth:               0 bps
  Average Total bandwidth:                0 bps

    By Type of Connection
                ---------
Unicast Incoming Bandwidth:               0 bps
  Multicast Incoming Bandwidth:              0 bps
Unicast Outgoing Bandwidth:               0 bps
  Multicast Outgoing Bandwidth:              0 bps

        By Type of Content
                    ------
    Live Incoming Bandwidth:              0 bps
     VOD Incoming Bandwidth:              0 bps
    Live Outgoing Bandwidth:              0 bps
     VOD Outgoing Bandwidth:              0 bps
```

```
            Overall Traffic
                   -------
          Incoming Bytes:               0 Bytes
          Outgoing Bytes:         2103939 Bytes
            Total Bytes:          2103939 Bytes

        Incoming Packets:                 0
        Outgoing Packets:              1403
           Total Packets:              1403

Movie Streamer Error Statistics
                                Total
---------------------------------------
        Server Error
            --------
     Internal Error:                          0
     Not Implemented:                         0
  Server Unavailable:                         0
     Gateway Timeout:                         0
              Others:                         0

        Client Error
            --------
         Bad Request:                         0
      File Not Found:                         0
   Session Not Found:                         0
  Method Not Allowed:                         0
Not Enough Bandwidth:                         0
    Client Forbidden:                         0
              Others:                         0

Movie Streamer Performance Statistics
                                  Total
-----------------------------------------
        CPU Usage:                     0.000000 %
           UpTime:                        5416 sec
Statistics were last cleared on Thursday, 25-Oct-2007 23:53:59 UTC.
```

Note In this example, current connections and bandwidth are all zero because the movie has finished playing at the time the statistics were displayed.

Note This example shows the Release 2.3 statistics, which were enhanced. All remaining examples for the Movie Streamer statistics show the Release 2.0-2.2 version.

Step 6 Play the movie again, and before it completes, display the statistics again.

As the statistics show, there are two current connections: 1 RTP and 1 RTSP.

```
NE_DEMO-SE2# show statistics movie-streamer all
Movie Streamer Request Statistics
                                Total
---------------------------------------
Current RTSP sessions:                  1
  Total RTSP sessions:                  1
   Current RIP connections:                 1
Total RTP connections:                  2

Movie Streamer Bandwidth Usage Statistics
                                  Total
```

```
---------------------------------------
Current Incoming bandwidth:               0 bps
Current Outgoing bandwidth:               0 bps
   Current Total bandwidth:               0 bps

Average Incoming bandwidth:               0 bps
Average Outgoing bandwidth:          863879 bps
  Average Total bandwidth:           863879 bps

         Incoming Bytes:                0 Bytes
         Outgoing Bytes:         14490660 Bytes
           Total Bytes:          14490660 Bytes
```

Testing Dynamically Ingested Movie Streamer Content

When content requested by a client is not in the CDS, it is dynamically ingested from the original source and streamed to the client by an SE.

Step 1 On the client PC, start the Apple QuickTime Player.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL of a sample MOV file that has not been preingested into the CDS in the Open URL dialog box (Figure G-15) and click **OK**. In the example, this is "rtsp://rfqdn.cds.com/test_dyn.mov."

Figure G-15 Open URL Dialog Box

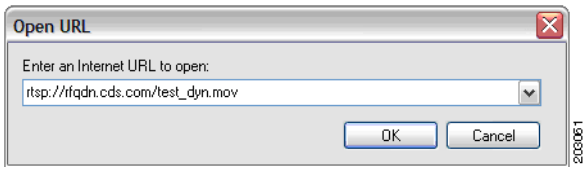

The video begins to play.

Step 4 To view the request flow, display the statistics. For CDS Releases 2.0 and 2.1, enter the **show statistics rtsp server movie-streamer all** command. For CDS Releases 2.2 and 2.3, enter the **show statistics movie-streamer all** command.

```
NE_DEMO-SE1# show statistics movie-streamer all
Movie Streamer Request Statistics
                              Total
---------------------------------------
Current RTSP sessions:                1
  Current RIP connections:                1
    Total RTP connections:                4

Movie Streamer Bandwidth Usage Statistics
                                 Total
-----------------------------------------
Current Incoming bandwidth:          207115 bps
Current Outgoing bandwidth:              0 bps
   Current Total bandwidth:          207115 bps

Average Incoming bandwidth:          199244 bps
Average Outgoing bandwidth:          207932 bps
```

```
  Average Total bandwidth:            407176 bps

          Incoming Bytes:           12667891 Bytes
          Outgoing Bytes:           12609164 Bytes
            Total Bytes:            25277055 Bytes

NE_DEMO-SE2# show statistics movie-streamer all
Movie Streamer Request Statistics
                               Total
---------------------------------------
Current RTSP sessions:                    0
   Current RIP connections:                   0
Total RTP connections:                    3

Movie Streamer Bandwidth Usage Statistics
                                  Total
-----------------------------------------
Current Incoming bandwidth:           194974 bps
Current Outgoing bandwidth:               0 bps
   Current Total bandwidth:           194974 bps

Average Incoming bandwidth:           174557 bps
Average Outgoing bandwidth:               0 bps
  Average Total bandwidth:            174557 bps

          Incoming Bytes:           13283705 Bytes
          Outgoing Bytes:                 0 Bytes
            Total Bytes:            13283705 Bytes
```

Testing Movie Streamer Live Content Playback

Note Each live program uses a live delivery service to deliver the live program. See the "Configuring Programs" section on page 5-26 for more information.

Live streaming of Movie Streamer content requires a Session Description Protocol (SDP) file. The SDP file used in this procedure is the following:

```
v=0
o=QTSS_Play_List 2079157989 2079176415 IN IP4 171.71.50.216
s=C:\Program Files\Darwin Streaming Server\Playlists\untitled\unti@
c=IN IP4 127.0.0.1
b=AS:94
t=0 0
a=x-broadcastcontrol:RTSP
m=video 0 RTP/AVP 96
b=AS:79
a=rtpmap:96 X-SV3V-ES/90000
a=control:trackID=1
m=audio 0 RTP/AVP 97
b=AS:14
a=rtpmap:97 X-QDM/22050/2
a=control:trackID=2
a=x-bufferdelay:4.97
```

Step 1 On the client PC, start the Apple QuickTime Player.

Step 2 Choose **File > Open URL**.

Step 3 Enter the URL of the live program file in the Open URL dialog box (Figure G-16) and click **OK**. In the example, this is "rtsp://rfqdn.cds.com/movieliveprogram.sdp."

Figure G-16 Open URL Dialog Box

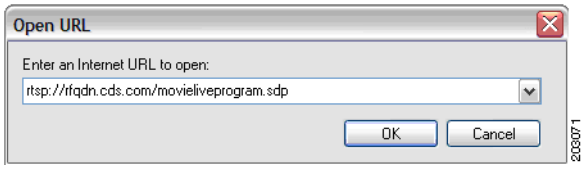

The video begins to play.

Step 4 To view the statistics on the video file, choose **Window > Show Movie Inspector** (Figure G-17).

The current statistics are displayed and updated as the video streams and plays.

Figure G-17 QuickTime Player Statistics

Step 5 To view the incoming and outgoing streams, display the statistics. For CDS Releases 2.0 and 2.1, enter the **show statistics rtsp server movie-streamer all** command. For CDS Releases 2.2 and 2.3, enter the **show statistics movie-streamer all** command.

```
NE_DEMO-SE1# show statistics movie-streamer all
Movie Streamer Request Statistics
                              Total
---------------------------------------
Current RTSP sessions:                  1
  Current RIP connections:                  1
Total RTP connections:                  3

Movie Streamer Bandwidth Usage Statistics
                                 Total
-----------------------------------------
Current Incoming bandwidth:          161526 bps
Current Outgoing bandwidth:          582640 bps
```

```
   Current Total bandwidth:           744166 bps

Average Incoming bandwidth:           192102 bps
Average Outgoing bandwidth:           203980 bps
  Average Total bandwidth:            396082 bps

          Incoming Bytes:            4478769 Bytes
          Outgoing Bytes:            4499370 Bytes
            Total Bytes:             8978139 Bytes

NE_DEMO-SE2# show statistics movie-streamer all
Movie Streamer Request Statistics
                                Total
--------------------------------------
Current RTSP sessions:                    0
  Current RIP connections:                    0
Total RTP connections:                    1

Movie Streamer Bandwidth Usage Statistics
                                   Total
-----------------------------------------
Current Incoming bandwidth:           175399 bps
Current Outgoing bandwidth:              0 bps
   Current Total bandwidth:           175399 bps

Average Incoming bandwidth:              0 bps
Average Outgoing bandwidth:              0 bps
  Average Total bandwidth:               0 bps

          Incoming Bytes:            1248165 Bytes
          Outgoing Bytes:            1080984 Bytes
            Total Bytes:             2329149 Bytes

Movie Streamer Performance Statistics
                                Total
--------------------------------------
       CPU Usage:                    0.000000 %
          UpTime:                     78375 sec
Statistics were last cleared on Friday, 26-Oct-2007 20:09:42 UTC.
```

Testing the Flash Media Streaming Engine

Note Flash Media Streaming is a Release 2.1 feature.

This section consists of the following procedures:

- Testing Flash Media Streaming Preingested Content
- Testing Flash Media Streaming Dynamically Ingested Content
- Testing Flash Media Streaming—Live Streaming

Testing Flash Media Streaming Preingested Content

Note Content must be preingested. This was accomplished in the "Configuring Delivery Services" section on page 5-1.

Note All RTMP calls must be in the following format:

```
rtmp://rfqdn/vod/path/foo.flv
```

In this format, *rfqdn* is the routing domain name of the Service Router, *vod* is the required directory, and *path* is the directory path to the content file that conforms to the standard URL specification.

Step 1 On the client PC, start the Adobe Flash Player.

Step 2 Enter the URL of the flash file in the text box (Figure G-18) and click **Go**. In the example, this is "rtmp://rfqdn.cds.com/vod/test_prepos.flv."

The RTMP call is routed to a Service Engine by the Service Router.

The FLV file has been preingested on the SEs. The video begins to play.

Figure G-18 Adobe Flash Player

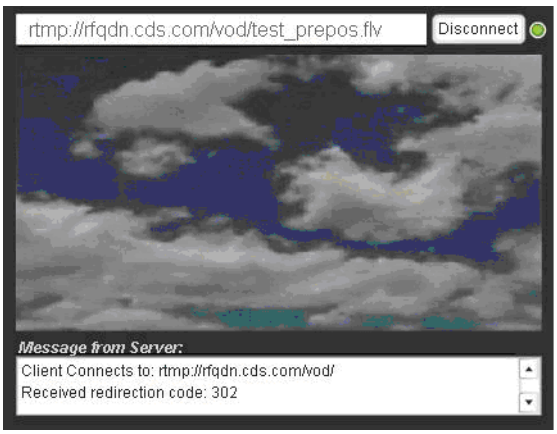

Step 3 To view the Flash Media Streaming statistics, enter the **show statistics flash-media-streaming** command on the SEs.

There is one concurrent connection on NE-DEMO-SE2, which means there is an active connection to this SE. The statistics also show a Preposition Hit of 1, which means there was preingested content being requested through this SE.

```
NE-DEMO-SE2# show statistics flash-media-streaming
Flash Media Streaming Statistics
Statistics were last cleared on Thursday. 06-Dec-2007 37:22:58 UTC.

Connections
-----------
Current         :          1
Current VOD     :          1
Current LIVE    :          0
Max Concurrent  :          1
Total           :          1
```

```
Live Streaming
-----
UpStream BW    :          0 kbps
DownStream BW  :          0 kbps
UpStream Bytes :          0
DownStream Bytes:         0
Num of Instance Load:     0

Performance
-----------
Server Up Time :          816 S
Mem Usage      :           5 %
Max Mem Usage  :           5 %

Cache
-----
Cache Hit      :          0
Cache Miss     :          0
Proxy Case     :          0
Cache HIt Percentage:    0.00 %

Preposition
-----------
Preposition Hit:          1

Bytes Served
------------
Local Disk Reads:         3
HTTP Based Reads:         0
Bytes From Local Disk: 880668
Bytes Through HTTP:       0

Rules
------------
Action Allow   :          0
Action Block   :          0
Validate url Sign:        0
Errors         :          0

Error
-----
Invalid Error  :          0
Server Error   :          0
Media Not Found:          0
Media Unauthorized:       0
Invalid Request:          0
```

Note The live streaming and rules statistics, as well as current VOD and current live, were introduced in Release 2.3

Testing Flash Media Streaming Dynamically Ingested Content

Note Testing dynamically ingested content for Flash Media Streaming includes cache miss scenarios and live streaming scenarios. Flash Media Streaming uses RTMP to stream live content by means of dynamic proxy. Configuration of live or rebroadcast programs is not required. When the first client requests live streaming content, the stream is created.

Step 1 SSH to NE-DEMO-SE1 and use the **show cache content** command to verify there is no cached content.

```
NE-DEMO-SE2# show cache content
Max-cached-entries is set as 3000000
Number of cal cached assets: 0
-------------------------------------------------
Priority    Size        URL
-------------------------------------------------
```

Step 2 On the client PC, start the Adobe Flash Player.

Step 3 Enter the URL of the flash file in the text box (Figure G-19) and click **Go**. In the example, this is "rtmp://rfqdn.cds.com/vod/test_cache.flv."

The RTMP call is routed to a Service Engine by the Service Router.

The FLV file has been cached on the SEs. The video begins to play.

Figure G-19 Flash Player

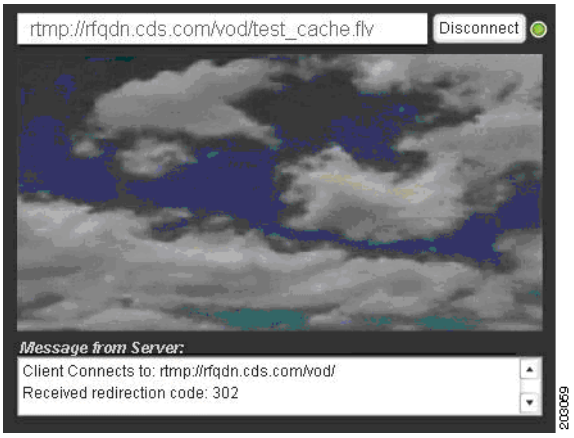

Step 4 To view the Flash Media Streaming statistics, enter the **show statistics flash-media-streaming** command on the SEs.

There is one concurrent connection on NE-DEMO-SE2, which means there is an active connection to this SE. The statistics also show a Cache Miss of 1, which means the content was not found on this SE.

```
NE-DEMO-SE2# show statistics flash-media-streaming
Flash Media Streaming Statistics
Statistics were last cleared on Thursday. 06-Dec-2007 37:22:58 UTC.

Connections
-----------
Current         :          1
Current VOD     :          1
Current LIVE    :          0
```

```
Max Concurrent :          1
Total          :          1

Live Streaming
-----
UpStream BW    :          0 kbps
DownStream BW  :          0 kbps
UpStream Bytes :          0
DownStream Bytes:         0
Num of Instance Load:     0

Performance
-----------
Server Up Time :          933 S
Mem Usage      :           5 %
Max Mem Usage  :           5 %

Cache
-----
Cache Hit      :          0
Cache Miss     :          1
Proxy Case     :          0
Cache HIt Percentage:    0.00 %

Preposition
-----------
Preposition Hit:          0

Bytes Served
------------
Local Disk Reads:         2
HTTP Based Reads:         1
Bytes From Local Disk: 587112
Bytes Through HTTP:    293556

Rules
------------
Action Allow   :          0
Action Block   :          0
Validate url Sign:        0
Errors         :          0

Error
-----
Invalid Error  :          0
Server Error   :          0
Media Not Found:          0
Media Unauthorized:       0
Invalid Request:          0
```

Note The live streaming and rules statistics, as well as current VOD and current live, were introduced in Release 2.3

Step 5 To verify that the content has been cached after it was requested, enter the **show cache content** command.

```
NE-DEMO-SE2# show cache content
Max-cached-entries is set as 3000000
Number of cal cached assets: 1
-------------------------------------------------
Priority   Size       URL
```

```
-------------------------------------------------
1.71425e-01293556     http://ofqdn.cds.com/vod/test_cache.flv
```

Testing Flash Media Streaming—Live Streaming

Note Live streaming using Flash Media Streaming is a Release 2.2 feature.

Flash Media Streaming uses RTMP to stream live content by means of dynamic proxy. Configuration of live or rebroadcast programs is not required. When the first client requests live streaming content, the stream is created. There are no limits to the number of live streams other than the system load. Live streaming uses distributed content routing to distribute streams across multiple Service Engines.

Step 1 Set up a Flash Media encoder. Enter the following information:

- FMS URL—Origin Server URL (Origin Server cannot be a CDS device.)
- Stream—Stream name for the client's request
- Video—Choose VP6 or H.264

Step 2 Click **Start** to publish the stream to the Origin Server.

Step 3 In a web browser on the client PC, enter the URL rtmp://<*edge SE IP address*>/live/<*publish stream name*>.

For example, if the URL was rtmp://Temp4.se.fmslive.com/live/livestream, *Temp4* is the SE assigned under the delivery service, the *live* directory indicates that it is a live stream, and *livestream* is the published name on the Flash Media Encoder.

Step 4 On the Edge SE enter the **show statistics flash-media-streaming** command to view the Flash Media Streaming statistics.

```
NE-DEMO-SE2# show statistics flash-media-streaming
Flash Media Streaming Statistics
Statistics were last cleared on Thursday. 06-Dec-2007 37:22:58 UTC.

Connections
-----------
Current        :          1
Current VOD    :          0
Current LIVE   :          1
Max Concurrent :          1
Total          :          3

Live Streaming
-----
UpStream BW    :          0 kbps
DownStream BW  :        274 kbps
UpStream Bytes :       3194
DownStream Bytes:    124362967
Num of Instance Load:     1
```

Step 5 Enter the **show flash-media-streaming livestreams** command.

```
NE-DEMO-SE2# show flash-media-streaming livestreams
Display flash-media-streaming livestreams
```

```
Display maximum 4 applications, 8 forwarders, and 8 client stream info

Forwarder
---------
ClientReqUrl       : rtmp://Temp4.se.fmslive.com/live/livestream
path2OS            : 2.225.2.65->2.225.2.62
reqFwdUrl          : rtmp://2.225.2.65/live/cds_fms_proxy/2.225.2.62/live/livestream
Upstream BW (Kbps) : 0
Downstream BW (Kbps): 267
Upstream Bytes     : 3487
Downstream Bytes   : 138278573
numClient          : 1
```

Step 6 On the Content Acquirer in the delivery service, enter the **show statistics flash-media-streaming** command. The command shows one session connecting from the edge SE.

```
NE-DEMO-SE2# show statistics flash-media-streaming
Flash Media Streaming Statistics
Statistics were last cleared on Thursday. 06-Dec-2007 37:22:58 UTC.

Connections
-----------
Current        :          1
Current VOD    :          0
Current LIVE   :          1
Max Concurrent :          1
Total          :          1

Live Streaming
-----
UpStream BW    :           0 kbps
DownStream BW  :         274 kbps
UpStream Bytes :        3487
DownStream Bytes:    149456451
Num of Instance Load:      1
```

Step 7 Enter the **show flash-media-streaming livestreams** command. The command shows the client request URL connecting from the edge SE.

```
NE-DEMO-SE2# show flash-media-streaming livestreams
Display flash-media-streaming livestreams

Display maximum 4 applications, 8 forwarders, and 8 client stream info

Forwarder
---------
ClientReqUrl       : rtmp://2.225.2.65/live/cds_fms_proxy/2.225.2.62/live/livestream
path2OS            : 2.225.2.62
reqFwdUrl          : rtmp://2.225.2.62/live/livestream
Upstream BW (Kbps) : 0
Downstream BW (Kbps): 261
Upstream Bytes     : 3337
Downstream Bytes   : 155736798
numClient          : 1
```

APPENDIX H

Software Licensing Information

This appendix provides software license information related to the Internet Streamer CDS.

Notices

The following notices pertain to this software license.

OpenSSL/Open SSL Project

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

License Issues

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License:

Copyright © 1998-2007 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)”.

4. The names “OpenSSL Toolkit” and “OpenSSL Project” must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

   “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)”.

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT “AS IS”' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License:

Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young’s, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

   “This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)”.

   The word ‘cryptographic’ can be left out if the routines from the library being used are not cryptography-related.

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: “This product includes software written by Tim Hudson (tjh@cryptsoft.com)”.

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License].